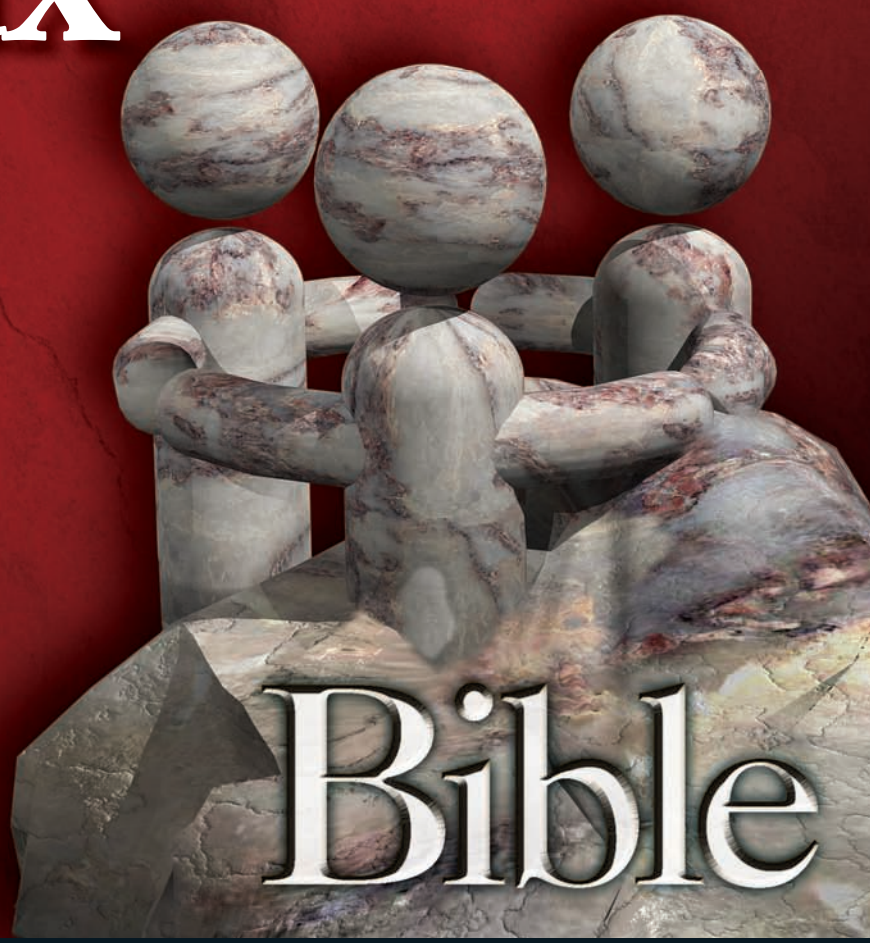


Complete Kubuntu® Coverage

- DVD includes Ubuntu 8.10 installation
- 14 bonus chapters included on CD-ROM

William von Hagen

Ubuntu® 8.10 Linux®



Bible

Set up Ubuntu for desktops or networks

Download music, watch DVDs, or play games

Harness easy-to-use administrative tools

The book you need to succeed!

Ubuntu® 8.10 Linux® Bible

Ubuntu® 8.10 Linux® Bible

William von Hagen



Wiley Publishing, Inc.

Ubuntu® 8.10 Linux® Bible

Published by
Wiley Publishing, Inc.
10475 Crosspoint Boulevard
Indianapolis, IN 46256
www.wiley.com

Copyright ©2009 by Wiley Publishing, Inc., Indianapolis, Indiana

Published simultaneously in Canada

ISBN: 978-0-470-29420-8

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at http://www.wiley.com/go/permissions.

Limit of Liability/Disclaimer of Warranty: The publisher and the author make no representations or warranties with respect to the accuracy or completeness of the contents of this work and specifically disclaim all warranties, including without limitation warranties of fitness for a particular purpose. No warranty may be created or extended by sales or promotional materials. The advice and strategies contained herein may not be suitable for every situation. This work is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional services. If professional assistance is required, the services of a competent professional person should be sought. Neither the publisher nor the author shall be liable for damages arising herefrom. The fact that an organization or Web site is referred to in this work as a citation and/or a potential source of further information does not mean that the author or the publisher endorses the information the organization or Web site may provide or recommendations it may make. Further, readers should be aware that Internet Web sites listed in this work may have changed or disappeared between when this work was written and when it is read.

For general information on our other products and services please contact our Customer Care Department within the United States at (877) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Library of Congress Cataloging-in-Publication Data

Von Hagen, William.
Ubuntu 8.10 Linux bible / William von Hagen.
p. cm.
Includes index.
ISBN 978-0-470-29420-8 (paper/cd-rom)
1. Ubuntu (Electronic resource) 2. Linux. 3. Operating systems (Computers) I. Title.
QA76.76.O63V664 2009
005.4'32--dc22

2009007444

Trademarks: Wiley and the Wiley logo are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates, in the United States and other countries, and may not be used without written permission. Ubuntu is a registered trademark of Canonical Limited. Linux is a registered trademark of Linus Torvalds. All other trademarks are the property of their respective owners. Wiley Publishing, Inc. is not associated with any product or vendor mentioned in this book.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

To Dorothy, for more than words can say . . .

Quidquid latine dictum sit, altum sonatur.

About the Author

William von Hagen (Bill) has been a UNIX system administrator for over twenty years, and a Linux fanatic since the early 1990s. He has worked as a Linux product manager, systems programmer, system administrator, writer, application developer, drummer, and content manager. Bill has written or co-written books on such topics as Linux server hacks, GCC, Xen virtualization, Linux filesystems, SUSE Linux, Red Hat Linux, SGML, Mac OS X, and hacking the TiVo. He has also written numerous articles on Linux, embedded computing, Mac OS X, UNIX, and various open source topics. An avid computer collector specializing in workstations, he owns more than 200 computer systems but is not compulsive at all. You can reach him at `vonhagen@vonhagen.org`.

Credits

Executive Editor
Carol Long

Senior Development Editor
Tom Dinse

Production Editor
Daniel Scribner

Copy Editors
Nancy Rapoport, Foxxe Editorial Services

Editorial Manager
Mary Beth Wakefield

Production Manager
Tim Tate

Vice President and Executive Group Publisher
Richard Swadley

Vice President and Executive Publisher
Barry Pruett

Associate Publisher
Jim Minatel

Project Coordinator, Cover
Lynsey Stanford

Compositor
Jeffrey Wilson, Happenstance Type-O-Rama

Proofreader
Kristy Eldredge
Scott Klemp
Amy Morales
Nate Pritts
Kyle Schlesinger

Indexer
Ron Strauss

Cover Image
Joyce Haughey

Cover Designer
Michael E. Trent

Acknowledgments

Nothing I have ever written would have been possible without the love, support, and infinite patience of Dorothy Fisher, my wife and best friend. Dorth believed in me from day one, removed several thousand commas from my earliest writing, and has accepted "I can't do that now—come back in fifteen minutes or six hours" as an excuse more times than anyone should ever be expected to. I am similarly lucky to have great friends like Jeff "Sunshine" Kaminski, Gregg Kostelich, Dr. Joe O'Lear, Jim Morgan, and Kim Walter in my life. I heard once that a good friend will come and bail you out of jail, but a true friend will be sitting in jail with you saying, "Man! That was fun!" Luckily, we haven't had to test that, but I'm sure we'd all look great in orange.

I would also like to thank Carol Long, Tom Dinse, Daniel Scribner, and others at Wiley who enabled me to do a new edition of this book and supported me during its creation. All errors are mine alone, but this would be a much weaker book without your contributions. I quite literally wouldn't have been able to do it without you. I would also like to thank the many readers who sent comments and suggestions after they had read the first edition of this book, especially Atanas Zheliazkov, Chris Cunha, and Reima Kangasniemi, who inspired me to devote equal time to Kubuntu.

Finally, this book wouldn't exist without Linus Torvalds, Mark Shuttleworth, the Ubuntu and Kubuntu folks in general, the Debian Project, Richard Stallman, the FSF, and the millions of contributors to the cornucopia that is GNU/Linux.

Contents at a Glance

Introduction xxxi

Part I: Getting Started With Ubuntu and Kubuntu 1
Chapter 1: The Ubuntu Linux Project 3
Chapter 2: Installing Ubuntu and Kubuntu 27
Chapter 3: Installing Ubuntu and Kubuntu on Special-Purpose Systems 65

Part II: Ubuntu and Kubuntu for Desktop Users. 101
Chapter 4: Basic Linux System Concepts 103
Chapter 5: Using the GNOME Desktop 123
Chapter 6: Using the Compiz Window Manager 175
Chapter 7: Using the KDE 3 Desktop 203
Chapter 8: Using the KDE 4 Desktop 273
Chapter 9: Managing E-mail and Personal Information with Evolution 313
Chapter 10: Managing E-mail and Personal Information with Kontact 371
Chapter 11: Surfing the Web with Firefox 421
Chapter 12: Migrating from Windows Systems 455
Chapter 13: Sending and Receiving Instant Messages 493
Chapter 14: Using Command-Line Tools 535
Chapter 15: Working with Text Files on Ubuntu 569
Chapter 16: Creating and Publishing Documents 605
Chapter 17: Other Office Software: Spreadsheets and Presentations 659
Chapter 18: Working with Graphics 705
Chapter 19: Working with Multimedia 745
Chapter 20: Would You Like to Play a Game? 825
Chapter 21: Consumer Electronics and Ubuntu/Kubuntu 859

Part III: Ubuntu and Kubuntu for System Administrators 903
Chapter 22: Understanding the System Startup and Shutdown Processes 905
Chapter 23: Adding, Removing, and Updating Software 925
Chapter 24: Adding Hardware and Attaching Peripherals 993
Chapter 25: Network Configuration and Security 1049
Chapter 26: Going Wireless 1091
Appendix A: What's on the DVD and CD-ROM? 1109

Index 1113

Part IV: Bonus Chapters (CD only) 1
Chapter B1: Connecting to Other Systems 3
Chapter B2: File Transfer and Sharing 29
Chapter B3: Software Development on Ubuntu and Kubuntu 65
Chapter B4: Managing Users, Groups, Authentication, and Advanced Permissions 109
Chapter B5: Managing System Policies and Security 151
Chapter B6: Backing Up and Restoring Files 171
Chapter B7: Using Virtual Machines and Emulators 203
Chapter B8: Setting Up a Web Server 277
Chapter B9: Setting Up a Mail Server 287
Chapter B10: Setting Up a DHCP Server 307
Chapter B11: Setting Up a DNS Server 321
Chapter B12: Setting Up a Print Server 339
Chapter B13: Setting Up an NFS Server 355
Chapter B14: Setting Up a Samba Server 371

Contents

Introduction xxxi

Part I: Getting Started With Ubuntu and Kubuntu 1

Chapter 1: The Ubuntu Linux Project 3

Background 4
Why Use Linux? 4
What Is a Linux Distribution? 5
Introducing Ubuntu Linux 6
The Ubuntu Manifesto 7
Ubuntu and Kubuntu Release Schedule 8
Ubuntu Update and Maintenance Commitments 9
Ubuntu and the Debian Project 10
Why Choose Ubuntu or Kubuntu? 11
What Versions of Ubuntu Are Available? 12
Installation Requirements 16
Supported System Types 17
Hardware Requirements 17
Time Requirements 18
Ubuntu CDs 18
Support for Ubuntu and Kubuntu 19
Community Support and Information 19
Documentation 22
Commercial Support for Ubuntu Linux 23
Getting More Information About Ubuntu 24
Summary 25

Chapter 2: Installing Ubuntu and Kubuntu 27

Ubuntu, Kubuntu, and Which Kubuntu 28
Getting a 64-Bit or Kubuntu Desktop CD 29
Booting from a Desktop CD 30
Installing Ubuntu or Kubuntu from a Desktop CD 34
Booting Ubuntu or Kubuntu 45
Booting Ubuntu or Kubuntu on Dual-Boot Systems 46
The First Time You Boot Ubuntu or Kubuntu 47
Test-Driving Ubuntu or Kubuntu Linux 47
Exploring the Desktop CD's Examples Folder 48

Accessing Your Hard Drive from the Desktop CD 50
Copying Files to Other Machines Over a Network 56
Using Desktop CD Persistence 56
Wubi: Installing Ubuntu or Kubuntu on Windows 60
Summary 63

Chapter 3: Installing Ubuntu and Kubuntu on Special-Purpose Systems 65

Getting a Different Install CD 66
Booting from a Server or Alternate Install CD 66
Install Options on the Server Install CD 67
Installing an Ubuntu Server 68
Manually Specifying Your Partition Layout 84
Booting Your Server for the First Time 93
Install Options on the Alternate Install CD 94
Installing a Desktop System in Text Mode 96
Installing in OEM Mode 97
Installing a Command-Line System 98
Installing an LTSP Server 99
Summary 99

Part II: Ubuntu and Kubuntu for Desktop Users 101

Chapter 4: Basic Linux System Concepts 103

Working with Files and Directories 103
Standard Linux Directories 105
Other Common Directories on Linux Systems 105
Introduction to Linux Filesystems 106
Disks, Partitions, and Mount Points 106
Introducing Logical Volumes 107
Local Filesystems: Standard and Journaling 109
Network Filesystems 110
Working with Partitions and Filesystems 111
Mounting Filesystems 111
Automatically Mounting Filesystems at Boot Time 114
Automatically Mounting Removable Media Filesystems 117
Understanding Linux Permissions 117
Basic Concepts: Users and Groups 118
File and Directory Permissions Under Linux 119
Performing Privileged Operations in Ubuntu 120
Summary 121

Chapter 5: Using the GNOME Desktop 123

What's a Desktop? Graphical Environments for Linux 123
Using the Mouse 126
GNOME Desktop Overview 126

GNOME Application Windows 128
Menus in GNOME 129
Panel Menus 130
Context-Sensitive Menus 132
Customizing GNOME Menus 134
Customizing Your Desktop 138
Customizing Mouse Behavior 138
Configuring Screen Resolution 140
Customizing Panels 142
Configuring the Screensaver 147
Configuring Power Management 148
Changing Desktop Backgrounds 152
Switching Themes 153
Assistive Technologies for Using GNOME 155
GNOME Keyboard Shortcuts 156
Introducing the Nautilus File Manager 157
Basic Operations in Nautilus 158
Examples of Using Nautilus 159
Getting More Information About Nautilus 162
Using the Tracker Desktop Search Application 162
Searching with Tracker 163
Configuring Tracker 165
Using the GNOME Deskbar Applet 167
Using a Window Manager 170
Summary 174

Chapter 6: Using the Compiz Window Manager 175

What Is Compositing? Why All the Fuss? 176
Compiz History: Compiz, Emerald, Beryl, and Friends 176
Enabling Special Effects in Compiz 178
Using the Appearance Preferences Application 178
Detailed Compiz Configuration 180
Other Useful Packages for Compiz 192
Installing and Using the Avant Window Navigator 192
Using Emerald and the Emerald Theme Manager 197
Using Compiz with KDE 200
Summary 201

Chapter 7: Using the KDE 3 Desktop 203

A Short Course in KDE History 204
Using the Mouse 205
KDE 3 Desktop Overview 206
KDE 3 Application Windows 208
Menus in KDE 3 212
Panel Menus 212

Context-Sensitive Menus 215
Customizing KDE Menus 217
Customizing the KDE 3 Desktop 223
Customizing Mouse Behavior 224
Configuring Screen Resolution 227
Customizing Panels 229
Configuring a Screensaver 242
Configuring Power Management 244
Changing Desktop Backgrounds 247
Switching Themes 248
Undoing KDE Configuration Changes 252
KDE Keyboard Shortcuts 252
Introducing the Dolphin File Manager 253
Overview of Dolphin 253
Basic Operations in Dolphin 257
Getting More Information About Dolphin 258
Using the Konqueror Browser 259
Configuring Konqueror 261
Using Konqueror as a File Manager 262
Using the Strigi Desktop Search Application 264
Changing the Window Manager for KDE 268
Installing KDE on an Ubuntu System 271
Summary 271

Chapter 8: Using the KDE 4 Desktop 273

What's New in KDE 4 274
KDE 4 Desktop Overview 277
Using the Mouse 277
KDE 4 Window Decorations 278
Menus in KDE 4 280
Running Commands with KRunner 288
Using the KDE 4 Panel 289
Using Widgets 290
Adding Widgets 291
Default Widgets in KDE 4 293
Getting Other Widgets 296
KWin and Compositing 297
Using Compositing in KWin 297
Standard Compositing Effects and Shortcuts 299
Configuring Advanced Graphics Options 302
Enabling and Disabling Compositing Effects 303
Customizing Your Desktop 308
Configuring Screen Resolution 309
Switching and Creating Themes 311
Summary 312

Chapter 9: Managing E-mail and Personal Information with Evolution 313

Starting Evolution 314
Using the Evolution Setup Assistant 315
Sending and Receiving Mail 326
Sending Mail in Evolution 330
Setting Evolution Preferences for Sending and Viewing Mail 331
Undeleting Mail Messages 334
Creating and Using Mail Folders 334
Using Search Folders 336
Filtering Incoming Mail 338
Automatically Processing Incoming Mail 338
Automatically Checking for Junk Mail 341
Accessing Free and Web-Based E-mail 342
Reading Gmail from a Desktop Mail Client 343
Reading Mail from Other Web-Based Mail Providers 347
Managing Contacts 349
Configuring Evolution's Contact Support 349
Adding and Editing Contact Information 354
Defining Mailing Lists 357
Managing Your Calendar 359
Basic Calendar Configuration 359
Creating and Editing Personal Calendar Entries 361
Creating and Editing Meetings 364
Integrating Online Calendars 366
Publishing Your Free/Busy Information 367
Additional Sources of Information About Evolution 369
Summary 370

Chapter 10: Managing E-mail and Personal Information with Kontact 371

Starting Kontact 371
Configuring Kontact 372
Configuring E-mail in Kontact 373
Configuring Contacts 386
Configuring Your Calendar 390
Configuring the Summary View 395
Sending, Receiving, and Storing Mail 396
Checking and Reading Mail in Kontact 396
Setting Preferences for Reading Mail in Kontact 398
Sending Mail in Kontact 404
Setting Preferences for Sending Mail in Kontact 406
Organizing and Filtering Mail 408
Using Folders to Organize Your Mail 408
Using Filters to Process Incoming Mail 409
Configuring Spam and Virus Detection in Kontact 411
Creating and Managing Contacts 414
Using and Managing Your Calendar 416
Summary 419

Chapter 11: Surfing the Web with Firefox 421

A Quick History of Firefox 421
Starting Firefox 422
The Firefox User Interface 424
Standard Parts of a Firefox Window 424
Using the Mouse in Firefox 426
Special and Not-So-Special Firefox Features 427
Configuring Firefox 429
Setting Your Home Page 429
Controlling Pop-ups, JavaScript, and More 430
Configuring Your Privacy Settings 431
Working with Bookmarks 434
Creating Bookmarks in Firefox 435
Managing Your Bookmarks 437
Creating Live Bookmarks 439
Enhancing Firefox 442
Registering for Add-ons 442
Viewing Installed Firefox Add-ons and Themes 442
Adding Firefox Add-ons and Themes 444
Adding New Firefox Plug-ins 450
Summary 453

Chapter 12: Migrating from Windows Systems 455

Overview of Migration Scenarios 456
Backing Up Your Personal Data 457
Migrating E-mail to Linux 458
Migrating Mail via Thunderbird 459
Manually Migrating Mail 473
Migrating Web/Browser Settings to Linux 475
Migrating from Firefox to Firefox 476
Migrating from Internet Explorer/Safari to Firefox 479
Migrating Just Your Bookmarks 481
Working with Windows Application Data 489
Summary 491

Chapter 13: Sending and Receiving Instant Messages 493

Instant Messaging Overview 494
Getting an Instant Messaging Account 496
Instant Messaging on Ubuntu with Pidgin 498
Configuring Instant Messaging Accounts in Pidgin 498
Using Pidgin 501
Getting More Information about Pidgin 509
Instant Messaging on Kubuntu with Kopete 510
Configuring Instant Messaging Accounts in Kopete 510
Using Kopete 513
Getting More Information about Kopete 525

Using IRC 525
Using IRC in Pidgin 526
Using IRC in Kopete 528
Getting More Information about IRC 531
Other Instant Messaging Packages 531
Resolving IM Communication Problems 532
Summary 533

Chapter 14: Using Command-Line Tools 535

Why Use the Command Line? 536
Executing Commands from the Command Line 536
What's a Shell? 542
Getting to a Shell 544
Using the GNOME Terminal Application 545
Using the KDE Konsole 546
Using the X Window System Terminal Application 548
Popular Command-Line Commands 551
Moving Around in the Linux Filesystem 552
Copying, Moving, Renaming, and Deleting Files and Directories 552
Changing File and Directory Permissions 556
Finding Commands That Do Specific Things 559
Working with the Bash Shell 559
Using Command History 560
Using Command and Filename Completion 561
Using Wildcards 562
Pipes and Input and Output Redirection 563
Introducing Job Control 564
Exploring the Bash Configuration File 565
Using Environment Variables 566
Defining and Using Aliases 568
Summary 568

Chapter 15: Working with Text Files on Ubuntu. 569

Introduction to Linux Text Editors 570
Using vi 571
Starting and Exiting vi 573
Inserting Text in vi 574
Moving Around in vi 574
Deleting and Changing Text in vi 575
Cutting, Copying, and Pasting in vi 576
Searching for and Replacing Text in vi 577
Undoing Changes in vi 578
Using Multiple Windows in vi 579
Customizing vim 580
Graphical Versions of vi 582
More Information About vi and vim 583

Using emacs 583
A Few Words About emacs Commands 586
Emacs Terminology 587
Starting and Exiting emacs 588
Moving Around in emacs 591
Cutting, Copying, and Pasting in emacs 592
Searching for and Replacing Text in emacs 592
Working with Multiple Windows and Buffers in emacs 592
Customizing emacs 593
Getting More Information About emacs 595
Using gedit 595
Popular KDE Text Editors 597
Using KEdit 597
Using KWrite 599
Using Kate 600
Other Text Editors for Linux Systems 602
Summary 603

Chapter 16: Creating and Publishing Documents 605

Using Document Markup Languages 606
Installing TeX and LaTeX 606
Using TeX Markup 607
Introducing LaTeX Markup 608
Formatting and Previewing DVI Files 610
Using Graphical Tools to Work with TeX and LaTeX 611
More Information About TeX and LaTeX 615
Word Processing with OpenOffice.org Writer 615
Installing Files for Writer 617
Taking a Quick Tour of Writer 618
Personalizing Writer 620
Using Wizards to Create Documents 622
Modifying Document Styles and Layout 625
Creating and Using Templates 627
Importing Documents from Other Word Processors 630
More Information About Writer 633
Word Processing (and More) with KWord 633
Installing KWord 634
Quick Tour of KWord 634
Creating and Editing Documents in KWord 637
Working with Frames 638
Importing, Exporting, and Saving Documents 645
More Information About KWord 645
Desktop Publishing with Scribus 645
Installing Scribus 646
Taking a Quick Scribus Tutorial 647
More Information About Scribus 656

Other Word Processors and Office Suites 657
Summary 657

Chapter 17: Other Office Software: Spreadsheets and Presentations 659

Introduction to Spreadsheets: A Quick Tutorial 660
Creating and Using OpenOffice.org Spreadsheets 669
Installing Files for OpenOffice.org Calc 670
Starting Calc 671
Taking a Quick Tour of Calc 672
Invoking Functions in Calc 674
Specifying the Type of Data in a Cell 674
Importing Existing Spreadsheets into Calc 677
Creating and Using Gnumeric Spreadsheets 680
Starting Gnumeric 681
Taking a Quick Tour of Gnumeric 681
Invoking Functions in Gnumeric 682
Specifying the Type of Data in a Cell 684
Importing Existing Spreadsheets into Gnumeric 686
Creating and Using KSpread Spreadsheets 688
Starting KSpread 688
Taking a Quick Tour of KSpread 689
Invoking Functions in KSpread 690
Specifying the Type of Data in a Cell 691
Importing Existing Spreadsheets into KSpread 693
Creating and Using OpenOffice.org Presentations 695
Installing Files for Impress 695
Starting Impress 696
Taking a Quick Tour of Impress 699
Creating a Presentation 700
Importing Existing Presentations 702
Summary 704

Chapter 18: Working with Graphics 705

Overview of Digital Graphics Terminology 706
Using GIMP 709
Starting GIMP 710
A Quick Tour of GIMP 711
Sample GIMP Tasks 714
More Information About GIMP 721
Using Krita 721
Starting Krita 723
A Quick Tour of Krita 724
Test-Drive: Using Krita 725
More Information About Krita 729
Photo Editing Overview 730
Using OpenOffice.org Draw 733
Starting Draw 733

A Quick Tour of Draw 733
More Information About OpenOffice.org Draw 735
Vector Graphics Tools for Linux 735
Using Inkscape 736
Using Karbon14 739
Other Graphics Applications 742
Summary 743

Chapter 19: Working with Multimedia 745

Overview of Digital Audio and Video Terminology 745
Configuring Sound Devices, Levels, and System Sounds on Ubuntu Systems 750
Specifying and Testing Sound Devices 751
Setting System Sound Levels 753
Testing and Customizing System Sounds 755
Getting Detailed Information about Your Sound Hardware 757
Setting Blank CD and DVD Preferences 758
Configuring Sound Devices, Levels, and System Sounds on Kubuntu Systems 760
Specifying System Sound Devices 760
Setting System Sound Levels 765
Testing and Customizing System Sounds 766
Getting Detailed Information about Your Sound Hardware 768
Setting Blank CD and DVD Preferences 769
Exploring the Linux Multimedia Stack 772
Introducing Drivers, Servers, and Frameworks 773
Configuring the GStreamer Framework and Plug-Ins 775
Configuring PulseAudio on Ubuntu 776
Playing CDs and DVDs 779
Playing Audio CDs Using Rhythmbox on Ubuntu 780
Playing Audio CDs with Amarok on Kubuntu 780
Installing Software to Play Encrypted DVDs 782
Playing DVDs with Totem on Ubuntu 783
Playing DVDs on Kubuntu Systems 784
Ripping Audio CDs 787
Ripping Audio CDs Using Rhythmbox on Ubuntu 787
Ripping Audio CDs Using k3b 789
Backing Up DVDs from the Command Line 792
Creating CDs and DVDs 793
Burning CDs Using Brasero on Ubuntu 793
Burning Data CDs and DVDs Using Nautilus 799
Burning CDs and DVDs Using k3b 801
Working with Other Audio Sources 808
Playing Audio Files and Internet Radio Using Rhythmbox 809
Playing Audio Files and Internet Radio on Kubuntu 814
Converting Audio File Formats 820
Summary 824

Chapter 20: Would You Like to Play a Game? 825

Default Games on Ubuntu 825
Card Games 826
Board Games 827
Video Games 829
Games on Kubuntu Systems 830
Arcade Games 831
Board Games 835
Card Games 837
Games for Kids 838
Tactics & Strategy 839
Games on the Games Menu 841
Other Popular Games 842
Billions and Billions of Versions of Chess 842
Go 845
Monopoly Clones 847
Scrabble Clones 851
Unique Arcade Games for Linux 852
Summary 857

Chapter 21: Consumer Electronics and Ubuntu/Kubuntu. 859

Configuring Your System for Consumer Electronics Devices 860
Digital Device Recognition on Ubuntu Systems 860
Digital Device Recognition on Kubuntu Systems 865
Working with Your Digital Camera 873
Working with Digital Cameras on Ubuntu Systems 873
Working with Digital Cameras on Kubuntu Systems 879
Working with iPods and Other Digital Audio Players 884
Digital Audio Players and Rhythmbox on Ubuntu 885
Digital Audio Players and Amarok on Kubuntu 888
Working with an iPod Using gtkpod 896
Working with CompactFlash and SD Cards 901
Summary 901

Part III: Ubuntu and Kubuntu for System Administrators 903

Chapter 22: Understanding the System Startup and Shutdown Processes . . 905

Overview of the Ubuntu/Kubuntu Startup Process 905
The Boot Monitor or BIOS 906
The Boot Loader 906
Loading the Kernel 909
Loading and Using an Initial RAM Disk or RAM Filesystem 909
The Init Process 911
Examining the Boot Process with Bootchart 916

Optimizing the Boot Process 917
Minimizing the Amount of Software You Start 918
Optimizing Application Loading 919
The Ubuntu/Kubuntu Shutdown Process 921
Summary 923

Chapter 23: Adding, Removing, and Updating Software 925

Installing Software with the Add/Remove Menu Item 927
Using the Add/Remove Menu Item in Ubuntu 927
Using the Add/Remove Programs Menu Item in Kubuntu 929
Package Management Software Overview 931
Ubuntu Repositories and Components 932
Enabling Additional Repository Components 934
Other Repositories of Interest 936
Enabling Additional Repository Sources Using the Software Sources Tool 936
Problems Adding or Accessing Nonstandard Repositories 942
Exploring Your System Using dpkg and Friends 942
Listing the Packages That Are Installed on Your System 943
Listing the Packages That Are Available for Your System 944
Listing Information About a Package 947
Listing the Contents of a Package 947
Determining What Package Provides an Existing File 948
Determining What Package Provides a Missing File 949
Using apt-get to Add and Remove Software 950
Upgrading Your System Using apt-get 951
Smart System Upgrades Using apt-get 951
Retrieving Package Source Code Using apt-get 952
Satisfying Build Dependencies Using apt-get 953
Using aptitude to Add and Remove Software 954
Tips and Tricks for Using the aptitude User Interface 955
Using aptitude to Install Recommended Software 956
Advantages of Using aptitude to Install and Remove Software 958
Using Synaptic to Add and Remove Software 959
Configuring Synaptic Preferences 961
Searching for Software in Synaptic 962
Installing Packages in Synaptic 964
Removing Packages in Synaptic 968
Using the Ubuntu Update Manager 970
Using the Adept Tools 972
Using the Adept Manager on KDE 3-Based Systems 972
Using the Adept Updater on KDE 3-Based Systems 980
Using the Adept Manager on KDE 4-Based Systems 980
Using the Adept Updater on KDE 4-Based Systems 987
Converting Packages from Other Package Formats 988
Keeping Your System Lean, Mean, and Pristine 991
Summary 992

Chapter 24: Adding Hardware and Attaching Peripherals 993
Adding a Printer 994
Adding a Printer on Ubuntu 995
Adding and Configuring Printers on Kubuntu 1002
Using a Scanner 1018
Using a Scanner on a Ubuntu System 1018
Using a Scanner on a KDE 3-Based System 1023
Adding Internal Disks and CD/DVD Drives 1027
Adding EIDE/ATA Drives 1027
Adding SATA Drives 1029
Adding SCSI Drives 1030
Troubleshooting Boot Problems after Adding New Drives 1030
Locating, Partitioning, and Formatting New Drives 1032
Using PCMCIA Cards 1043
Adding PCI Cards 1044
Examining and Troubleshooting Devices Graphically 1045
Summary 1047

Chapter 25: Network Configuration and Security 1049
Networking 101 1050
Configuring Your Network Hardware 1053
Using Network Manager on Ubuntu 1054
Using the Network Manager on Kubuntu 1067
Configuration Files for Network Interfaces 1077
Network Testing with Network Tools 1080
Tips for Securing Your System 1082
Installing a Firewall 1083
Overview of Linux Firewalling and Packet Filtering 1084
Installing and Configuring a Firewall 1085
Summary 1090

Chapter 26: Going Wireless 1091
Overview of Wireless Technologies 1091
Using Wicd 1094
Installing Wicd 1095
Using Wicd 1096
Configuring Wicd 1100
Command-Line Tools for Wireless Networking 1102
Installing and Using Windows Networking Drivers 1104
Installing NDIS Wrapper and Friends 1104
Installing Microsoft Windows Drivers 1105
NDIS Wrapper Tips and Tricks 1107
Summary 1108

Appendix: What's on the DVD and CD-ROM? 1109
System Requirements 1109
Using the DVD with Windows 1110
Using the DVD with Linux 1110
What's on the DVD 1110
What's on the CD 1111
Troubleshooting 1112
Customer Care 1112

Index 1113

Part IV: Bonus Chapters (CD only) 1

Chapter B1: Connecting to Other Systems 3
Establishing Secure Connections to Other Systems 3
Logging in to Remote Systems Using ssh 4
Enabling Incoming SSH Connections to Your System 5
Connecting to Other Systems Using VNC 7
Configuring a VNC Server 7
VNC Client and Server Software for Other Systems 12
Using Ubuntu's Remote Desktop Viewer 13
Using Kubuntu's Remote Desktop Connection 16
Connecting to Remote Windows Terminal Servers 20
Using rdesktop for Terminal Server Connections 20
Using Ubuntu's tsclient for Terminal Server Connections 22
Using Kubuntu's krdc for Terminal Server Connections 24
Using Other Remote Connection Software 27
Summary 27

Chapter B2: File Transfer and Sharing 29
Transferring Files to Other Systems 30
Using FTP 32
Accessing Shares on Windows Systems 45
Accessing NFS Directories from Linux Systems 53
Peer-to-Peer File Sharing on Ubuntu 55
Installing and Using gtk-gnutella 56
Using BitTorrent 58
Summary 63

Chapter B3: Software Development on Ubuntu and Kubuntu 65
Installing and Using Development Software 66
Installing GCC Compilers, make, and Friends 66
GCC Compiler Overview 66
Using GCC's C Compiler 67
Creating Simple Makefiles 68

Building Open Source Software 70
Building a Standalone Package 71
Building Add-On Software 78
Building Supported Software in Different Ways 79
Using an Integrated Development Environment 80
Popular IDEs for Linux 81
Installing Eclipse 82
Using Eclipse for C Application Development 83
Getting More Information about Eclipse 99
Source Code Control Software 99
The Concurrent Versioning System (CVS) 100
Installing and Using Subversion 100
Summary 108

Chapter B4: Managing Users, Groups, Authentication, and Advanced Permissions 109

Creating and Managing Users and Groups 110
Managing Users and Groups on Ubuntu Systems 110
Managing Users and Groups on KDE 3 Systems 116
Managing Users and Groups on KDE 4 Systems 121
PAMs and the Linux Authentication Process 125
PAM Configuration Files for Applications and Services 126
Example: PAMs Used by the Login Process 127
Configuration Files for Various PAMs 130
What If PAM Configuration Files Are Missing? 131
Customizing the sudo Command 132
Using ACLs for Sophisticated File Sharing 135
Overview of Linux ACLs 136
Installing ACL Commands 136
Activating Filesystem Support for ACLs 137
Working with ACLs from the Command Line 137
Working with ACLs Using Graphical Tools 140
Configuring Kerberos Authentication 142
Installing Kerberos Client and Server Commands 142
Configuring NTP on Servers and Clients 143
Setting Up a Kerberos Server 145
Setting Up a Kerberos Client 147
Summary 149

Chapter B5: Managing System Policies and Security 151

Overview and Comparison of AppArmor and SELinux 152
Identifying Resources 153
Rule Readability and Ease of Creation 154
Comparison Summary 154
Configuring and Using AppArmor 154
Installing and Configuring AppArmor 155

AppArmor Profiles and Profile Syntax 157
Additional Information about AppArmor 167
Configuring and Using SELinux 167
Installing SELinux 168
SELinux Tools 169
Additional Information about SELinux 170
Summary 170

Chapter B6: Backing Up and Restoring Files. 171

Backups 101 172
Why Do Backups? 172
Different Types of Backups 174
Verifying and Testing Backups 176
Deciding What to Back Up 177
Backup Software for Linux 177
Local Backup and Restore Software for Linux 179
Network-Oriented Backup Software for Linux 180
Backing Up Files to Local, Removable Media 182
Archiving and Restoring Files Using tar 182
Making an Up-to-Date Copy of a Local Directory Using cp 184
Making an Up-to-Date Copy of a Remote Directory Using rsync 185
Installing and Using BackupPC 187
Installing and Configuring BackupPC 188
Identifying Hosts to Back Up 190
Defining a Backup Using rsyncd 191
Defining a Backup Using SMB 193
Starting Backups in BackupPC 194
Restoring from Backups in BackupPC 195
Linux Time Machines 198
Summary 201

Chapter B7: Using Virtual Machines and Emulators 203

Overview of Virtualization and Emulation 204
Why Virtualize? 204
Types of Virtual Machines 205
What Kinds of Virtualization Can I Use? 207
Selecting a Virtualization Technology 208
Popular Emulation Software for Linux 208
Managing Virtual and Emulated Machines 209
Getting More Information About Virtualization 210
Installing and Using VirtualBox 211
VirtualBox Overview 211
Installing VirtualBox 212
Creating a Virtual Machine in VirtualBox 213
Installing an Operating System in a Virtual Machine 220
Installing VirtualBox Guest Additions 224

Starting and Stopping Virtual Machines 224
Networking and VirtualBox 227
Fixing Virtual Machine Boot/Installation Problems 233
Installing and Using VMware Server 234
Downloading and Installing VMware Server 235
Starting VMware Server 242
Creating a Virtual Machine 244
Starting a VMware Virtual Machine 249
Installing and Using KVM 252
Installing KVM and Related Packages 253
Creating a Disk for a KVM Virtual Machine 254
Installing an Operating System in a Virtual Machine 255
Starting a KVM Virtual Machine 257
Networking for KVM 257
Installing and Using Xen 261
Overview of Xen 261
Installing Xen 263
Creating Xen DomainU Configuration Files 263
Building and Booting a Paravirtualized DomainU 267
Creating a Hardware Virtual Machine DomainU 271
Getting More Information About Xen 272
Installing and Using QEMU 273
Summary 275

Chapter B8: Setting Up a Web Server 277

World Wide Web 101 277
Introduction to Web Servers and Apache 279
Installing Apache 280
Apache 2 File Locations 280
Configuring Apache 282
Troubleshooting 285
More Information 286
Summary 286

Chapter B9: Setting Up a Mail Server 287

Introduction to Mail Servers 288
Popular Linux Mail Transfer Agents 288
Why Run Your Own Mail Server? 290
Installing Postfix and Friends 291
Packages to Install 291
Built-in Postfix Configuration Models 292
Configuring Postfix 294
Postfix Configuration Files 295
Identifying Trusted Hosts and Domains 297
Rewriting Addresses in Outgoing Mail 297
Accepting Mail for an Entire Domain 298

Activating Qpopper for POP/POP3 Support 298
Adding Spam Filtering and Virus Scanning to Postfix 300
Greylisting via postgrey 300
Adding Postfix Parameters to Reject Bogus Mail 301
Integrating MailScanner, SpamAssassin, and ClamAV with Postfix 303
More Information 305
Summary 306

Chapter B10: Setting Up a DHCP Server 307

Overview of DHCP 308
Installing a DHCP Server 309
Managing a DHCP Server from the Command Line 311
Creating DHCP Configuration Files Using a Text Editor 311
Specifying Additional DHCP Server Configuration File Entries 313
Managing a DHCP Server Graphically 314
Troubleshooting DHCP 319
Summary 320

Chapter B11: Setting Up a DNS Server 321

Overview of DNS and BIND 322
Installing a DNS Server 325
Overview of BIND Configuration Files 325
Creating DNS Zone and Reverse Lookup Files 327
Using Common Entries for Zone and Reverse Lookup Files: SOA and $TTL 327
Creating Zone Files 329
Creating Reverse Lookup Files 331
Incorporating Zone and Reverse Lookup Files with BIND 332
Restarting and Testing Your Name Server 333
Troubleshooting DNS 335
Getting More Information About DNS and BIND 337
Summary 337

Chapter B12: Setting Up a Print Server 339

Linux and UNIX Printing History 339
Enabling Remote Hosts to Access Your CUPS Print Server 340
Enabling Remote Printing on Ubuntu and KDE 4-Based Systems 341
Enabling Remote Printing on KDE 3–Based Systems 344
Integrating Windows Printing with an Ubuntu Print Server 346
Integrating Mac OS X Printing with an Ubuntu Print Server 347
Troubleshooting Remote Printing 350
Checking the CUPS Log Files 351
Accessing Controls for Portions of the CUPS Web Interface 351
Handling Preformatted Print Jobs 352
Getting More Information About CUPS 352
Summary 353

Chapter B13: Setting Up an NFS Server 355

Overview of the Network File System 356
How NFS Works 357
Comparing Different Versions of NFS 359
Installing an NFS Server 360
Configuring an NFS Server 360
Configuring Common NFS Options 361
Configuring NFS Kernel Server Options 362
Setting Up the Pseudo-Filesystem and Exported Directories 363
Defining Exported Directories in /etc/exports 363
Mounting NFS Directories on Clients 365
Debugging NFS Problems 367
Verifying RPC Services 367
Verifying Filesystem Export 368
Debugging NFS Operations 369
Getting More Information About NFS 370
Summary 370

Chapter B14: Setting Up a Samba Server 371

Overview of Microsoft Windows File Sharing 371
Introducing Samba 373
Installing the Samba Server and Friends 373
Samba Server Configuration Essentials 374
Identifying Your Workgroup or Domain 374
Configuring Samba Authentication 376
Sharing Printers and Home Directories Using Samba 376
Verifying the Samba Configuration File 378
Testing Samba Availability and Services 379
Sharing Other Directories via Samba 379
Getting More Information About Samba 380
Summary 380

Introduction

A Linux distribution is basically the sum of the things that you need to run Linux on your computer. There are many different Linux distributions, each with its own target audience, set of features, administrative tools, and fan club, the latter of which is more properly known as a user community. Putting aside the downright fanatics, most of the members of the user community for any Linux distribution are people who just happen to find themselves using a distribution for one reason or another. These reasons range from what they've heard from friends, what CD or DVD came with a Linux magazine they bought, or which Linux book they happened to buy.

Ubuntu Linux and siblings such as Kubuntu Linux are the most exciting Linux distributions in years. Ironically, while Ubuntu itself is indeed new, it also comes with a respectable Linux pedigree. Ubuntu has direct roots in one of the oldest and best-known Linux distributions available, the Debian GNU/Linux distribution. The folks who initially created and supported Ubuntu, Canonical Ltd., started out as Debian fans who wanted a faster-moving, more up-to-date distribution than Debian provided. So, in the spirit of Linux and the open source movement, they made their own distribution, Ubuntu Linux, by incorporating the best of Debian, other Linux distributions, and open source applications, and added their own special sauce. Ubuntu lovers who wanted to use a different graphical user interface than the default one provided with Ubuntu then created Kubuntu using the same philosophy.

Ubuntu means "humanity to others" in Zulu. For the people who use and bring you Ubuntu Linux, this is not just a name with touchy-feely overtones. The special sauce in Ubuntu is a social and business commitment to Ubuntu users everywhere. Ubuntu releases occur regularly, every six months, and support and updates for any Ubuntu release are available for a minimum of eighteen months after that. More about that it in the first chapter, where you'll read more about Ubuntu, its philosophy, its community, and why the sum of those makes Ubuntu and Kubuntu different than any other Linux distribution.

In a nutshell, Ubuntu is a Linux distribution for people. While reading this book, you'll see that there are plenty of excellent technical reasons for using Ubuntu, even if you're a hardcore Linux propeller head. However, that's not the point of Ubuntu—Ubuntu is for people who want to use their computers and need a solid software foundation for doing so. Whether your focus is on writing code or surfing the Web, sending and receiving electronic mail, working with your digital photographs, watching DVDs, listening to music, and so on, Ubuntu offers the software that you need to do what you want to do.

Like any Linux distribution, you can freely download and install Ubuntu, but it gets even better. If you don't have access to a CD burner or simply don't have the time, the Ubuntu folks will send you CDs that you can either use to install Ubuntu or to run Ubuntu on your current computer system, without changing anything. That's more than free—it's revolutionary!

And finally, Kubuntu. The first edition of this book focused on Ubuntu Linux, which was fine but slighted the huge and exciting community of Ubuntu users who prefer the KDE desktop environment. At the expense of some increase in page count and Web-only content, this book now covers both Ubuntu and Kubuntu—not just mentioning KDE as an afterthought, but discussing both KDE 3-based and KDE 4-based Kubuntu systems as full and equal partners in the Ubuntu experience.

Who Should Read This Book

If you're reading this in a bookstore and are unsure about which Linux distribution to get started with, or whether to use Linux at all, this book is for you. Ubuntu and Kubuntu are complete, visually friendly, and community-oriented distributions that make it easy for you to get started using Linux. Ubuntu and Kubuntu are designed to be distributions for users, but as you'll see throughout this book, you can do anything that you want with them—from running your desktop to running servers and network services for the enterprise—thanks to the inherent power of Linux. Both Ubuntu and Kubuntu come with a tremendous selection of up-to-date software, and plenty more is quickly downloaded and installed thanks to their easy-to-use administrative tools. Ubuntu and Kubuntu are frequently updated, and there is no such thing as a Linux virus or "accidentally installed spyware." Linux is inherently secure.

The rich Ubuntu and Kubuntu user communities are a big win for new and existing users. There are places to ask questions and actually get answers. People seem to want to help. The Ubuntu and Kubuntu forums and mailing lists are live, constant demonstrations of the philosophical and social aspects of Ubuntu, which are discussed in Chapter 1. In a nutshell, every Linux distribution has fans, devotees, forums, and mailing lists, but Ubuntu's and Kubuntu's are the most exciting, usable, and useful that I've ever seen in my years of using Linux. You can feel the excitement. Ubuntu and Kubuntu bring the promise of internationalization alive—you can get versions of Ubuntu and Kubuntu for many languages and character sets, and more are actively on the way.

If you're already using Ubuntu or Kubuntu, this book should be equally useful to you because it explains how to use the standard applications provided with these distributions, how to do common system configuration and system administration tasks, and so on. If you're already using Ubuntu or Kubuntu on your desktop and want to do more with it, this book clearly explains how to install common servers for file-sharing, electronic mail handling, Web servers, and much more.

How This Book Is Organized

Ubuntu 8.10 Linux Bible is organized into four parts.

Part I: Getting Started with Ubuntu and Kubuntu

The first part of this book provides background information about Ubuntu and Kubuntu Linux, including instructions on installing either of them on your machine or simply taking them for a test drive from a Live CD. Chapter 1 introduces Ubuntu and the Ubuntu philosophy, and explores the Ubuntu and Kubuntu communities and the various Web sites where you can get information and assistance in using and configuring your system (if you need help *now*, rather than simply reading this book). Chapter 1 also discusses some of the main reasons why Ubuntu and Kubuntu are the right Linux distributions to use, and helps you choose between them. Chapter 2 explains how to install Ubuntu as your only operating system or as an alternate operating system on an existing computer system if you can't live without whatever you're already using. Chapter 3 explains how to experiment with Ubuntu even if you don't have a computer system or disk space to spare, and also discusses alternate CDs that you can download to install Ubuntu or Kubuntu on systems that don't meet the standard installation requirements, on systems that you are reselling as an OEM, or as a server operating system.

Part II: Ubuntu and Kubuntu for Desktop Users

Part II explores the rich set of applications that are available for Ubuntu and Kubuntu systems, and explains how to use them to accomplish the kinds of things that people use modern personal computers for—reading and sending e-mail; surfing the Web; creating documents and spreadsheets; playing games; and playing and managing audio CDs, online music files, and DVD movies. Chapter 4 provides a basic discussion of Linux and UNIX fundamentals if you are curious about how Linux systems are organized or about using a command line. Chapters 5 and 6 give you an overview of the graphical user interface provided by Ubuntu and explore its window manager in detail. Chapters 7 and 8 provide an overview of the graphical user interface provided by Kubuntu systems of different vintages. If you're migrating to Ubuntu or Kubuntu from a Microsoft Windows system, Chapter 12 explains how to move all of your existing e-mail, Internet Explorer Favorites, and other data to your new (and superior) operating system and applications. The rest of the chapters in Part II are task-oriented discussions of common tasks, such as editing files, reading and sending mail, surfing the Web, sending instant messages, printing things, working with multimedia, creating spreadsheets and presentations, sharing and transferring files to other computer systems that are on your local network or anywhere on the Internet, working with your iPod or other MP3 player, and much more. This part concludes with a discussion of application development on Ubuntu, highlighting standard GNU/Linux development tools such as compilers, automatic build tools, graphical Integrated Development Environments, and source code control systems—all of which are free, of course.

Part III: Ubuntu and Kubuntu for System Administrators

Anyone who sets up their computer to connect to the network, enable users to log in, and so on is an official junior system administrator as far as I'm concerned. This part explains how to do all of the system administration tasks that any home user will want to do, but also explores all of the system administration tasks that you or any other administrator will need to do when using Ubuntu or Kubuntu in a business setting. Chapter 22 begins by explaining what happens when your Ubuntu or Kubuntu boots or is shut down, and how you can monitor and expedite the boot process. Chapter 23 tells you how to keep your system up-to-date and how to locate and install new software, while Chapter 24 explains how to add new hardware to your Ubuntu system if you run out of space, need faster graphics, or need anything that causes you to upgrade your machine. Chapter 25 discusses network security and how to protect yourself from crackers and other local or Internet vandals. Chapter 26 explains how to use Ubuntu and Kubuntu on laptops, other wireless systems, and with common wireless devices such as keyboards, mice, and even headphones.

Part IV: Bonus Chapters

Books can only be so big, but there's a lot to learn about Ubuntu and Kubuntu systems. The second CD that is included with this book contains PDF versions of 14 additional chapters that will help you master your Ubuntu or Kubuntu system. The topics discussed in these chapters (which are listed in the table of contents of this book) include how to connect to remote systems graphically from your Ubuntu or Kubuntu system, how to transfer and share files to and from your system (including how to use peer-to-peer clients such as Bittorrent and others), how to develop software on your system, and how to use all of the most popular virtualization software that is available for your Ubuntu or Kubuntu system. Other chapters discuss how to set up a Web server, select and configure a mail server, use your Ubuntu or Kubuntu system as a centralized print server, provide basic file sharing and print services to Microsoft Windows and Mac OS X systems on your network, and how to set up servers for popular acronyms such as DNS, DHCP, and NFS.

There are many different organizational and typographical features throughout this book designed to help you get the most out of the information.

Conventions Used in This Book

Tips, Notes, and Cautions

Whenever the authors want to bring something important to your attention, the information will appear in a Caution, Tip, or Note.



This information is important and is set off in a separate paragraph with a special icon. Cautions provide information about things to watch out for, whether simply inconvenient or potentially hazardous to your data or systems.



Tips generally are used to provide information that can make your work easier—special shortcuts or methods for doing something easier than the norm.



Notes provide additional, ancillary information that is helpful, but somewhat outside of the current presentation of information.

Part I

Getting Started With Ubuntu and Kubuntu

IN THIS PART

Chapter 1
The Ubuntu Linux Project

Chapter 2
Installing Ubuntu and Kubuntu

Chapter 3
Installing Ubuntu and Kubuntu on Special-Purpose Systems

Chapter 1

The Ubuntu Linux Project

Personal computers and their operating systems have come a long way since the late 1970s, when the first home computer hit the market. At that time, you could only toggle in a program by flipping switches on the front of the machine, and the machine could then run that program and only that program until you manually loaded another, at which time the first program was kicked off the system. Today's personal computers provide powerful graphics and a rich user interface that make it easy to select and run a wide variety of software concurrently.

The first home computer users were a community of interested people who just wanted to do something with these early machines. They formed computer clubs and published newsletters to share their interests and knowledge — and often the software that they wrote for and used on their machines. Sensing opportunities and a growing market, thousands of computer companies sprang up to write and sell specific applications for the computer systems of the day. This software ranged from applications such as word processors, spreadsheets, and games to operating systems that made it easier to manage, load, and execute different programs.

Although the power and capabilities of today's personal computers is light-years beyond the capabilities of those early machines, the idea of writing software and freely sharing it with others never went away. While it never got much press because nobody was making money from it, free software (and often its source code) has continued to be available from computer clubs, bulletin boards systems, and computer networks such as today's Internet. The free software movement finally blossomed with three seminal events:

- The creation of the GNU Project (`www.gnu.org`) by Richard Stallman in 1983, a project dedicated to developing software whose source code would always be freely available
- The announcement of the Free Software Foundation (FSF) (`www.fsf.org`), initially dedicated to fundraising for the GNU project
- The introduction of a free operating system project in 1991 that came to be known as Linux, by a Finnish computer software student named Linus Torvalds

IN THIS CHAPTER

Introducing Ubuntu Linux

Choosing Ubuntu

Reviewing hardware and software requirements

Using Ubuntu CDs

Getting help with Ubuntu Linux

Getting more information about Ubuntu

The book that you hold in your hands wouldn't exist without these three events, the resulting shockwave of independence and empowerment, and the perpetuation of the community spirit throughout the computer industry that these events (and many related ones) caused. Sometimes, if you're lucky, the more things change, the more they stay the same.

The operating system and applications discussed in this book are free, and their source code is freely available. Anyone who wants to can build, install, and run them. A huge online community of users has sprung up around them, including specialized groups who create easily installed sets of this software, known as Linux distributions. This chapter explores the philosophy, community, and history behind one of the newest, and arguably the best, of these easily obtained, easily installed, and easy-to-use free software environments, known as the Ubuntu Linux distribution.

Background

The emergence of Linux, a freely available operating system, is a landmark event in modern personal computing. Today, Linux is arguably the most popular operating system in use for server environments, and is quickly gaining significant numbers of users as a personal computer operating system for home use. The following two sections provide some background on Linux if you're just considering adopting it or are unfamiliar with some standard Linux terms such as the idea of a Linux distribution. If you're already familiar with Linux and are interested in Ubuntu as your Linux distribution of choice, you can skip this section and go directly to "Introducing Ubuntu Linux."

Why Use Linux?

Presumably, you've bought this book because you want to use Linux, but if you're just reading to find out why you might want to do so, some common reasons for using Linux are the following:

- **Powerful, modern design:** Linux was designed from the ground up to enable you to run multiple programs at the same time and to provide services that your computer and others can use. Most other desktop computer operating systems, such as Microsoft Windows, started out as small operating systems that could run only one program at a time, and they have been trying to catch up ever since.
- **Freely available source code means no lock-in to a single vendor:** Regardless of the operating system that you're currently using, you may have encountered problems with applications or the way things worked. However, if you're using an off-the-shelf operating system from Microsoft or Apple, you can get fixes and updates only from Microsoft or Apple. Linux is open and free, so if you don't like the way that Red Hat's Linux works, how much it costs, or the type of customer support that's available, you can always switch to Novell's SUSE Linux, Mandriva Linux, or (preferably) to Ubuntu Linux.
- **Thousands of free, powerful applications:** Need a word processor? Download and install OpenOffice Writer, AbiWord, Kwrite, or dozens of others. Need a database? Download and install MySQL, PostgreSQL, or many others. Need to create graphics or manipulate digital photographs? It doesn't get much better than GIMP (GNU Image Manipulation Program). If anything, a problem with Linux can be that you have too many choices, none of which cost money.
- **Support for standards:** Linux and Linux applications are designed to support standards, because standards are the language of free intellectual commerce. Linux applications support modern application and data formats for audio, multimedia, document formatting, spreadsheet data, and

many more. Because Linux is open and free, there can be no such thing as a proprietary Linux data or application format. This not only fosters data exchange between Linux applications, but also guarantees that you'll always be able to get to your data.

- **Lower total cost of ownership:** If you want to use Linux on your desktop or throughout your business, it's free to obtain and there are legions of Linux wizards available who can help you do whatever you want with it. There are no licensing fees — if you need to pay for something, you can pay for updates and support from the vendor of your Linux distribution.
- **Stable, powerful, and virus-free:** Linux is a mature, multiuser system that is dependable, stable, has built-in security and is immune to viruses except through system administration slipups.

It used to be the case that using Linux required some amount of special knowledge, but that's basically not the case any longer. Linux distributions such as Ubuntu make Linux easy — or, more properly, they make Linux invisible. As you'll see throughout the rest of this book, Ubuntu Linux provides an easy-to-use operating system and all of the applications that you need to do almost anything. The goal of this book is to explore Ubuntu Linux and its sister distribution, Kubuntu, explain how to have fun and get work done with them, and to provide any special details, insights, or knowledge that you might need. Shhh! Your grandmother doesn't have to know that she's running Linux.

What Is a Linux Distribution?

If you've been curious about Linux for a while, you've probably noticed that a bewildering number of different versions of it seem to be available. Computer magazines and Linux-related Web sites discuss Red Hat Linux, SUSE Linux, the Novell Linux Desktop, Fedora Core Linux, and many other things ending in Linux, each available from a different company or organization. Understanding exactly what people mean when they say "Linux" is the key to understanding how so many different versions of the same thing can be available, but that requires a little insight into how personal computers actually work in terms of software.

When you install an operating system such as Linux, Microsoft Windows, or Apple's Mac OS X on your computer system, you're installing some amount of software that is invisible to any regular user because it runs behind the scenes. This software handles scheduling, starting, and stopping different programs, communicating with your computer's hardware, handling communications with peripherals such as your printer, and so on. This is generally known as system software because regular users don't directly interact with it, but it needs to be present and running to provide the services that application software relies on. The core piece of this system software is usually referred to as a *kernel* because it is the central piece of the operating system and everything else builds upon the fundamental services that it provides.

By itself, a kernel isn't very interesting — people don't actually want to run a kernel, they want to run applications. These applications depend on services that are provided both by the kernel and by other system software. For example, if you want to print a file, whatever application you're using needs to create a version of your file that is formatted in a way that your printer understands, and then schedule that file for printing. Another piece of software handles sending the formatted file to the printer, making sure that the file prints correctly, and so on.

In popular usage, "Linux" is the collective name for an operating system kernel and its associated applications. In reality, Linux is technically the name of just the kernel — most of the applications that anyone uses with Linux come from other free software projects. A *Linux distribution* is the correct term for a Linux kernel, a set of applications that can run on top of it (regardless of where they come from), and a tool to install everything and configure your system. Each company or organization that provides a Linux

distribution is taking advantage of the open source nature of the Linux kernel and the applications that run on top of it by putting together the "right" version of the Linux kernel with what they view as the "right" collection of core applications that anyone would want to run on top of it.

NOTE **Because many of the key applications that systems running the Linux kernel depend upon have their roots in the GNU project, the historically proper way of referring to a Linux distribution is as a GNU/Linux distribution. However, given the number of other projects that have made huge contributions to today's Linux distributions, this book simply refers to the term as Linux distributions, rather than as GNU/GNOME/KDE/TeX/*your-favorite-project-here*/ Linux distributions. This in no way minimizes the fundamental and huge contributions that the GNU projects and the FSF have made to modern computing. I'm an FSF member, and strongly suggest that you should be one, too. See `www.fsf.org/associate` for details. It's always a good idea to support the things that you believe in (and depend on).**

Of course, getting a CD or DVD that just contained a bunch of software would be next to useless without some easy way of installing it, configuring it so that it works with your particular computer system (identifying peripherals, setting it up to communicate over your network or with your ISP, creating user accounts, and so on). Therefore, anyone who puts together a Linux distribution also provides a tool for installing and configuring the system, which is generally what runs when you boot from a Linux CD or DVD for the first time. This installation and configuration tool generally leverages a package management system that makes it easy to add or remove sets of related applications, identifying dependencies between different software components to ensure that the applications that you install will actually execute correctly.

Linux distributions are the key to understanding how Linux can be free and sold at the same time. The source code for the Linux kernel and open source applications is indeed freely available from thousands of sites on the Internet. Anyone who wants it can get it, but putting it all together in an easily installable, usable form is another thing entirely. When people sell a Linux distribution, they are basically just charging you for the media that it comes on, the time and effort that they invested in putting it all together, and (in some cases) "charging in advance" for any customer support that you might need if you encounter installation or initial configuration problems.

Developing Linux distributions and making them widely available has been critical to the adoption of Linux as an operating system because these distributions have made it possible for people to actually install and use Linux, the GNU utilities, and so on.

Introducing Ubuntu Linux

Ubuntu Linux is a Linux distribution founded in 2004. Originally focused on the needs of desktop and laptop users, Ubuntu has branched out since then, and now also offers distributions focused on the needs of commercial users with its Ubuntu Server distribution, Ubuntu JeOS for virtualization platforms, and Ubuntu Mobile for mobile and embedded devices such as smart phones, Internet tablets, and so on. All of these flavors of Ubuntu Linux are products of the Ubuntu project sponsored by Canonical, Ltd. (`www.canonical.com`), a company founded by Mark Shuttleworth, a successful South African entrepreneur, long-time Debian Linux developer, and general open source advocate. Ubuntu is a Debian-based Linux distribution (more about that later in this chapter) that uses a graphical user interface known as GNOME as its desktop environment. (GNOME is discussed in detail in Chapter 5, "Using the GNOME Desktop.") Sister projects (officially known as "Ubuntu Editions") to Desktop Ubuntu include:

- Kubuntu, a version of Ubuntu that uses the KDE desktop environment instead of GNOME. Versions of KDE are discussed in Chapter 7, "Using the KDE 3 Desktop" and Chapter 8, "Using the KDE 4 Desktop," and KDE applications are discussed throughout this book.
- Xubuntu, a version of Ubuntu that uses the lighter-weight Xfce desktop.
- Gobuntu, a version of Ubuntu that uses completely open source software.
- Edubuntu, a version of Ubuntu that focuses on educational applications and popularizing the use of Linux in schools.

NOTE **All of the versions of Ubuntu are built and distributed in the same way, and simply target different groups of users. Aside from their target audience, the only real differences between them are how they are installed, the set of applications that they provide when they are first installed, and how the Linux operating system itself is pre-configured for each. For more information about these different editions of Ubuntu, and what differentiates them from each other, see the section later in this chapter entitled "What Versions of Ubuntu Are Available?"**

Everything has to have a name, but what is the Ubuntu in *Ubuntu Linux*? Not too surprisingly, the Ubuntu Linux Web site puts it best:

> Ubuntu is an ancient African word, meaning "humanity to others." Ubuntu also means "I am what I am because of who we all are." The Ubuntu Linux distribution brings the spirit of Ubuntu to the software world.

Although that may be a bit touchy-feely for some, it's hard to argue with success and commitment. In 2005, its first year of availability, Ubuntu Linux received awards such as the *Linux Journal*'s Reader's Choice award, *Tux Magazine*'s Reader's Choice 2005 for Favorite Linux Distribution award, Ars Technica's Best Distribution award, the UK Linux & Open Source Industry's Best Distribution award, and the Linux World Expo's Best Debian Derivative Distribution award. Not too shabby for the new distribution on the block. Since 2005, Ubuntu has almost always been the most popular Linux distribution listed on the central Linux distribution site `www.distrowatch.org`. Ubuntu's popularity with users extends to computer vendors as well — computer hardware vendors such as Dell Computer have selected Ubuntu as their distribution of choice to ship with new hardware.

Aside from its technical excellence and usability (and some good funding thanks to Mark Shuttleworth), much of the success to date of Ubuntu Linux grows from the fact that its creators and proponents are not just the traditional Linux fanatics, but are genuinely committed to creating and promoting a usable and easily managed Linux distribution for end users all over the world.

The Ubuntu Manifesto

The Ubuntu Manifesto is a mission statement phrased in the classic manifesto form much beloved of artistic and political movements. The Ubuntu Manifesto is available online in the Philosophy section of the Ubuntu Web site (`www.ubuntulinux.org/ubuntu/philosophy`). Its core ideas are as follows:

- Every computer user should have the freedom to run, copy, distribute, study, share, change, and improve his software for any purpose, without paying licensing fees.
- Every computer user should be able to use her software in the language of her choice.
- Every computer user should be given every opportunity to use software, even if the user works under a disability.

The first bullet is largely a clear restatement of the goals of open source software in general, but the second and third bullets are two of the big drivers for the success of Ubuntu.

Internationalization is the term for producing software that is capable of displaying all prompts, dialogs, system messages, and so on, in any user's native language and any specific character set used with that language. The term internationalization is such a mouthful that it is frequently referred to as *i18n* because the word *internationalization* consists of the letter "i" followed by 18 letters and ends with an "n." The two aspects of i18n are *translation*, ensuring that versions of operating system and application messages and text are available in other languages, and *localization*, which ensure that messages and text can be displayed in a language's native character set(s). Amusingly, localization is often referred to as *l10n*.

Linux and its applications have been focused on i18n for years, thanks to initiatives such as the Linux Internationalization Initiative (`www.li18nux.net/`, known as *Li18nux*) and the Free Standards Group's Open Internationalization Initiative (`www.openi18n.org/`). These initiatives focus on ensuring that open source applications take i18n into account when developing, maintaining, and enhancing code. Many of the structural enhancements to the last few releases of desktop environments such as GNOME and KDE have been related to making sure that these environments and their applications support different languages and character sets.

The key to successful internationalization is two-fold — not only do applications and graphical environments need to support multiple languages and character sets, but the translations of prompts, dialogs, and system messages have to be available. Ubuntu's focus on a truly usable Linux distribution for an international audience has helped it become a hub for translation and localization work in Linux (`www.ubuntulinux.org/community/participate#l10n`) and GNOME (its primary graphical environment — more about that later) through an online translation system known as Rosetta (`https://launchpad.net/rosetta`), documentation translation efforts, active mailing lists, and other resources.



For additional information about Linux Internationalization, see additional Web resources such as the i18nGurus' Linux Internationalization Resources page at `www.i18ngurus.com/docs/984813514.html` and the Linux Internationalization HOWTO, available online at sites such as `http://home.no.net/david/i18n.php`.

Although many of the structural enhancements to the last few releases of desktop environments such as GNOME and KDE have been related to internationalization, a great deal of work has also been done to make Linux graphical environments easier to use by people with disabilities. Ensuring that graphical applications provide keyboard or gesture shortcuts for all menu commands and dialog interaction has been a growing focus for GNOME, KDE, and graphical application development. Ubuntu's emphasis on usability is a boon to all computer users, regardless of whether or not they have a physical disability.



For additional information about Linux Accessibility projects and usability awareness, see additional Web resources such as the Linux Accessibility HOWTO (`www.tldp.org/HOWTO/Accessibility-HOWTO/`), the Linux Developers Accessibility HOWTO (`http://larswiki.atrc.utoronto.ca/wiki`), and the Linux Accessibility Resource Site (`http://larswiki.atrc.utoronto.ca/wiki`).

Ubuntu and Kubuntu Release Schedule

Given the pace of open source software development, it's important for anyone who depends on a Linux distribution to be able to get the latest and greatest kernels and versions of software packages. Kernel and associated device driver improvements provide security fixes, facilitate the use of the latest hardware, and often provide performance improvements in the handling of existing devices and protocols. The latest versions of software packages typically provide improvements in both capabilities and usability. Because

the open source community model virtually guarantees that thousands of improvements are in progress at any given moment, delivering an integrated and tested version of the latest and greatest Linux kernel and supported software packages is a complex task, but is one that is extremely important to the success and widespread adoption of any Linux distribution.

The Ubuntu folks deliver a fresh Ubuntu release every six months. These regular releases provide an up-to-date and tested kernel and a well-tested, integrated set of user software including the X Window system release from X.Org, the latest stable GNOME desktop, and core Linux and GNOME applications including Ubuntu-specific applications and customizations. Kubuntu releases, which use the same set of packages as Ubuntu with the exception of the KDE desktop and KDE-based applications, follow the same schedule as Ubuntu releases.

A regular release schedule is something that is unique in the Linux space, and has led to a unique approach to version numbering for Ubuntu releases. Traditional software releases are numbered according to major and minor release numbers, where the major release number is essentially arbitrary and generally indicates some major upgrade in functionality. For example, in traditional release numbering, version 4.2 is the minor release of version 4 of the software that follows version 4.1. Ubuntu uses major release numbers that identify the year in which the software was released, and what appear to be the minor numbers actually represent the month in which the release was made. Therefore, version 8.04 is the Ubuntu release from the fourth month of 2008. In typical, lighthearted Linux fashion, each Ubuntu release also has a nickname. The following are the Ubuntu releases to date:

- 4.10: Warty Warthog (October, 2004)
- 5.04: Hoary Hedgehog (April, 2005)
- 5.10: Breezy Badger (October, 2005)
- 6.06 LTS: Dapper Drake (June, 2006) a Long Term Support (LTS) release
- 6.10: Edgy Eft (October, 2006)
- 7.04: Feisty Fawn (April, 2007)
- 7.10: Gutsy Gibbon
- 8.04 LTS: Hardy Heron, a Long Term Support (LTS) release
- 8.10: Intrepid Ibex
- 9.04: Jaunty Jackalope (in beta at the time of this writing)

This is a refreshing approach to version numbering, especially in the open source space where software is essentially under continuous development. Not only does it make it possible for users and system administrators to predict and schedule system updates to their systems, but it also makes it easy to identify the vintage of existing, installed systems without consulting a reference text or the Web.

Ubuntu Update and Maintenance Commitments

The frequency with which a distribution is released is important to any user in order to guarantee that they have access to the latest and greatest system and application software. However, for any business that is interested in the power, cost savings, and flexibility of Linux, the period of time in which a release will be updated and maintained is even more important. Many businesses have hundreds or thousands of computer systems. Businesses with substantial infrastructure on top of deployed systems need to spend significant time testing updates and new releases before they can even think about rolling them out to all

of their computer systems. Once testing is complete, the physical act of updating deployed systems takes significant time, which literally translates into money in terms of MIS and IT personnel.

By default, security updates for each Ubuntu release are issued for 18 months after the release date — after that date, existing security updates are still available, but there is no guarantee that new updates for that release will be provided. To address the longer-term requirements of commercial Ubuntu users, Mark Shuttleworth and Canonical, Ltd., the founders and sponsors of Ubuntu Linux, formed and funded the Ubuntu Foundation in mid-2005. The Ubuntu Foundation employs some of the core Ubuntu community members to help guarantee the success and continuity of the Ubuntu development and release process. One aspect of the Ubuntu Foundation is an increased commitment to maintenance and updates. Ubuntu version 6.06 was the first Ubuntu Linux release to benefit from this when the Ubuntu Foundation announced that Ubuntu version 6.06 would be supported for three years on the desktop and five years on the server, doubling the standard 18-month support commitment that is a backbone of Ubuntu. The acronym LTS, for Long Term Support, was added to its release number to highlight that fact. Ubuntu 8.04 is the latest long-term support release. Long-term support commitments are valuable to any Ubuntu user because they guarantee that both support and updates will be available for an extended period of time, but are especially important for business users who need to select a stable and long-lived platform to satisfy their enterprise computing requirements.

Given the pace of hardware and open source software development, Ubuntu's regular release schedule and long maintenance commitment help ensure that Ubuntu users always have stable, secure, and up-to-date versions of the software that they depend on. Ubuntu also provides a firm update and support commitment that is mandatory for its successful commercial adoption of Ubuntu.

Ubuntu and the Debian Project

Debian is one of the longest-lived Linux distributions available, and is the Linux distribution that is the conceptual parent of Ubuntu Linux. Debian is pronounced Deb'-ian, with a soft e, and is a contraction of the names of the founders of the Debian Project, Debra and Ian Murdock.

The Debian Project (`www.debian.org`) was founded in 1993 and has been delivering quality GNU/Linux distributions ever since. Debian is well known for stable releases based on a huge collection of thoroughly tested and completely integrated software packages. Unfortunately, the downside of balancing testing and integration with keeping up with the pace of open source development has led to a painfully slow release history — there have literally been years between official Debian releases. To be fair, three versions of the current Debian release are always available: stable (the released version), testing (the candidate for the next release), and unstable (the development version). However, many businesses (and users) are uncomfortable depending on something labeled *testing* or *unstable*. Providing a faster release process, focusing on specific core technologies such as the GNOME desktop interface, and providing a better structured mechanism to deliver software updates and notify users of their availability are the key reasons why the Ubuntu project was born.

Some of the key ways in which the Debian and Ubuntu Linux distributions interact are the following:

- Ubuntu shares the software packaging format used by Debian (known as DEB packages), and also relies on the excellent and impressive technologies that were pioneered by Debian for identifying and resolving dependencies and relationships between different open source software packages.
- Ubuntu developers feed their changes and enhancements to open source packages back to the open source community, but also provide them directly to the Debian developers responsible for that package and even record patch information directly into the Debian bug-tracking system. Bug fixes and related enhancements made by Ubuntu developers are delivered as they are made

during the Ubuntu release and testing process, not in a big bang fashion once an Ubuntu release is complete. This is better for everyone.

- The Debian and Ubuntu distributions are based on a slightly different selection of open source packages, but follow the same general organization of those packages into separate domains, as explained in Chapter 23, "Adding, Removing, and Updating Software."

The Debian and Ubuntu Linux distributions are closely linked, complementary distributions with different goals. Ubuntu would not exist without the pioneering efforts and contributions of the Debian distribution, but provides a more predictable distribution with better support channels for many users and enterprise computing environments.

Why Choose Ubuntu or Kubuntu?

As mentioned earlier in this chapter, and in any recent computer magazine, and if you've ever looked at the Linux section in your local bookstore, there are zillions of different Linux distributions. After all, it's free, so why not? Techies aside, most of the users of any Linux distribution are people who have heard that they should be using a specific Linux distribution, got a free Linux CD for some distribution in a Linux magazine that they bought, or happened to buy a book about a specific Linux distribution.

Ubuntu means "humanity to others," but the title bar on the Ubuntu Web pages has often stated "Linux for People," and that's what Ubuntu is really about — a Linux distribution for people who want to get work done with a minimum of fuss and bother. Never mind that it's also a technically sophisticated Linux distribution with up-to-date software. Does it do what I want it to do?

The answer is unquestionably "Yes!" However, if you're unconvinced or find yourself in a discussion about Linux at a cocktail party, you may want more empirical data. Here are a few of the attributes of Ubuntu Linux that make it an attractive distribution to just about anyone:

- **Regular, up-to-date releases:** The Linux kernel and the thousands of software packages that make up the Linux user and administrative environment are constantly being updated. As discussed earlier in this chapter, providing the latest and greatest kernel and application software on a regular schedule is a fundamental principle of Ubuntu Linux.
- **Commitment to quality:** The quality of a Linux distribution hinges on two things: how good it is in the first place, and the distributing vendor's degree of commitment to fixing problems that arise. In both cases, Ubuntu shines. Each release goes through extensive internal testing by the Ubuntu team and extensive public testing of release candidates. Once a release occurs, updates for that release are delivered for a minimum of eighteen months (as needed, of course).
- **Community and commercial support:** Much of the support for any Linux distribution comes from its user community, and it's hard to beat the passion and commitment of the Ubuntu community. However, just as no business can afford to depend on an operating system without a reasonable maintenance commitment, no business can afford to depend on an operating system without some chance of guaranteed support. As discussed later in this chapter, a complete spectrum of commercial and community support is readily available for Ubuntu Linux.
- **Easy retrieval and application of updates:** The previous bullets have stressed the importance of being able to keep installed Linux systems up-to-date. Ubuntu provides great tools that notify users when updates are available and makes them easy to obtain and install. Ubuntu's graphical Update Manager and Synaptic Package Manager tools (discussed in Chapter 23, "Adding, Removing, and Updating Software") are the best examples of such tools that I have ever used.

- **Focus on usability:** Ubuntu defines itself as "Linux for People" and provides custom graphics, window decorations, and color schemes designed to provide an attractive, usable desktop environment for real people for personal use and to get work done. Like any other Linux distribution, you can customize this extensively, even switching to any of a variety of other window managers or desktop environments that are easily retrieved and installed through the Synaptic Package Manager. Ubuntu uses the GNOME desktop environment by default, which is well known for its support of and sensitivity toward accessibility requirements such as keyboard equivalents for menus and menu commands. Kubuntu uses the KDE desktop environment, which has always been regarded as a very user-friendly and easy-to-use desktop environment.
- **Focus on internationalization:** Ubuntu is extremely focused on supporting translation efforts and providing a Linux distribution that people anywhere on the planet can use in their native language, with their native character sets.
- **Active and involved community:** As I'll discuss in the next section, it is hard to conceive of a more active, dynamic, and involved user community than that which surrounds Ubuntu. An active and involved community translates into more places to ask questions, a better chance of getting answers, and a more friendly experience when doing so.

As you can see from this list, Ubuntu and Kubuntu focus on solving many of the issues that plague other Linux distributions or that make it difficult for new users to adopt Linux as their operating system of choice. Most general-purpose Linux distributions would claim that they address the same sorts of issues, but in my experience, Ubuntu and Kubuntu are exceptional in terms of delivering on them.

What Versions of Ubuntu Are Available?

As discussed earlier in this chapter, a number of different versions of Ubuntu are available directly from the folks at Canonical or from the distribution sites for sister projects. In addition to Ubuntu's "official" sister projects, a number of other Linux distributions are available that are based on Ubuntu, taking advantage of the rich set of packages that are supported by Canonical and the Ubuntu build infrastructure and distribution tools, just as Ubuntu itself leveraged the Debian Linux distribution to create the first Ubuntu distributions. The free and open source nature of most Linux code and applications makes it easy (well, relatively easy) to create different Linux distributions that are tailored to the needs and personal tastes of different types of users.

The standard Ubuntu Desktop distribution, using GNOME as its desktop environment, is my personal favorite for desktops and modern laptops, and is the one that I recommend for most people. It's hard to argue with success. On the other hand, other versions of Ubuntu may be better for you depending on your goals and personal tastes. This section discusses the different Ubuntu and Ubuntu-based distributions to help you decide which of these best suits your tastes and requirements, dividing these distributions based on the audience and type of user that they target.

NOTE **This book focuses on the GNOME-based Ubuntu Desktop, Kubuntu Desktop, and Ubuntu Server distributions. If you decide to use one of the other Ubuntu or Ubuntu-based distributions that are covered in the next few sections, the discussions of using various desktop applications, server applications, and the GNOME and KDE desktops will still be valuable to you. Unfortunately, there is no discussion of the applications and tools that are specific to other distributions. Finishing this book in my lifetime was a requirement from the publisher.**

Desktop and Laptop Users

By desktop and laptop computer users, I mean end users (like myself) who are typically looking for a Linux distribution that is stable and easy to maintain, and that provides a rich set of available applications. End users typically want to have some fun and get some work done. This isn't limited to home computing, and includes office, development, and educational environments. The following are the Ubuntu and Ubuntu-based distributions that I am most familiar with that target end users. You'll see what differentiates them from other flavors of Ubuntu, and where to get them:

- **Edubuntu: A**n official edition of Ubuntu that focuses on educational applications and popularizing the use of Linux in schools. Edubuntu uses GNOME as its default desktop environment, and provides a number of different desktop themes that are designed to be attractive (and useful) for users of different age groups. Edubuntu provides an incredibly rich set of educational applications, such as the KDE Edutainment suite (education and entertainment, get it?), GCompris for kindergarten users, Tux4Kids (paint, math, and typing applications), and the OpenOffice suite of desktop applications. You can find out more about Edubuntu or download a copy from the Edubuntu web site at `www.edubuntu.org`. A Live CD version of Edubuntu is available if you want to experiment with it before permanently installing an Ubuntu distribution.
- **Fluxbuntu:** An extremely lightweight Linux distribution that is based on Ubuntu, and which targets machines with limited amounts of memory, relatively small disks, and so on. It refers to itself as an LPAE (Lightweight, Productive, Agile, Efficient) Linux distribution, which is both accurate and a fine addition to your collection of computer-related abbreviations. Fluxbuntu's official memory requirements are less than 64MB, and its disk requirements are under 1.5GB, making it ideal for older or inexpensive machines that you don't want to invest in additional hardware for. Fluxbuntu isn't limited to older or low-power hardware — for example, a 64-bit version of Fluxbuntu is available. Fluxbuntu uses the Fluxbox window manager as its default graphical environment, and includes AbiWord as its primary word-processing application, GnuMeric as its default spreadsheet application, and lightweight Web applications such as Kazehakase (web browser, `http://kazehakase.sourceforge.jp/`), Claws (e-mail client, `www.claws-mail.org/`), and Pidgin (instant messaging, discussed later in this book). Like any Ubuntu-based distribution, many other packages are available. You can find out more about Fluxbuntu or download a copy from the Fluxbuntu web site at `www.fluxbuntu.org`.
- **Geubuntu:** Another lightweight Linux distribution that is based on Ubuntu. Unlike Fluxbuntu, Geubuntu uses E17, which is the window manager formerly known as Enlightenment (`www.enlightenment.org/`), which now includes various desktop features and the Xcompmgr compositing manager. Enlightenment has always been known as one of the most attractive and configurable window managers, and this Ubuntu-based distribution is an excellent choice for use on lower-power machines or on systems for which you prefer to devote as much of your CPU and memory as possible to the processes you're running versus the desktop environment. The founder of the Geubuntu effort is an Italian designer, so the aesthetics of Geubuntu are attractive and compelling. Geubuntu is available both as a live/install CD, and via a script that enables you to add its packages to an existing Ubuntu installation. For more information about Geubuntu or to download a copy, see the Geubuntu web site at `http://geubuntu.intilinux.com/Home.html`.
- **Gobuntu:** An official edition of Ubuntu that uses completely open source software with no restrictions on use or redistribution. The primary goal of Gobuntu is to provide a powerful and usable Ubuntu-based distribution that adheres to the Free Software Foundation's four freedoms (`www.fsf.org/licensing/essays/free-sw.html`). Gobuntu therefore does not include any firmware, drivers, applications, or other content for which the complete source code is not available and whose license does not provide the right to use, study, modify, and redistribute the body

of work. Gobuntu is the most philosophically pure Ubuntu-based distribution available. However, because of the restrictive licensing terms that are greedily imposed on device drivers and much audio, video, and other multimedia software, you may find it hard to use Gobuntu on certain hardware or to play certain types of audio or video files on Gobuntu. If this isn't critical to you and you are more philosophically pure than I am, you can find out more about Gobuntu or download a copy from the Ubuntu site at `www.ubuntu.com/products/whatisubuntu/gobuntu`.

- **Kubuntu:** An official edition of Ubuntu that uses the KDE desktop environment instead of GNOME. The KDE desktop environment is a powerful, extremely usable desktop environment for modern UNIX and UNIX-like computer systems that pre-dates the GNOME desktop and has a large and fanatical following for good reason. Many people prefer KDE to GNOME, and the discussions/arguments about selecting one desktop interface over the other have consumed just as much time and disk space as any other type of religious argument. Kubuntu combines the traditional ease-of-use of KDE with the stability and power of Ubuntu. Old and new Linux distributions such as Ark, Corel, Linspire, Mandrake, Mandriva, Mepis, OpenSUSE, Sabayon, SUSE, and Xandros use KDE as their default desktop environment. If you are already a user of one of these distributions or simply like what you've heard about KDE, Kubuntu may be the distribution for you. For more information about the KDE desktop itself, see Chapters 7 and 8, aptly titled "Using the KDE 3 Desktop" and "Using the KDE 4 Desktop," respectively. This book also includes discussions of all of the KDE-based equivalents to standard Ubuntu applications. For more information about KDE, see the official KDE web site at `www.kde.org`. You can find out more about Kubuntu or download a copy from their web site at `www.kubuntu.org`. A Live CD version of Kubuntu is available if you want to experiment with it before permanently installing an Ubuntu distribution.
- **Ubuntu Christian Edition:** A standard version of Ubuntu that has been augmented by adding Christian software such as content filtering packages (DansGuardian, pre-configured to filter many offensive web sites), daily Bible verse software, GnomeSword (a Bible study tool), BibleMemorizer, e-Sword and The Word (two truly impressive Bible study and research tools), Firefox with Christian themes and the WhatWouldJesusDownload toolbar (one of my favorite package names ever!), and much more. Beyond simply addressing some of the requirements of Christian computer users, this distribution also demonstrates the tremendous range of available software for Linux and how easy it is to customize Ubuntu to address the requirements of a specific audience. This distribution is currently available only for the 32-bit x86 systems. For more information about Ubuntu Christian Edition or to download a copy, see the web site at `www.whatwouldjesusdownload.com/christianubuntu`.
- **Ubuntu Muslim Edition:** A standard version of Ubuntu that has been augmented by adding Islamic software such as Islamiccal (an Islamic calendar), content-filtering packages (DansGuardian and Tinyproxy, preconfigured to filter many offensive web sites), Minibar (an application that identifies daily prayer times), and Zekr (a Qur'an study tool). Beyond simply addressing some of the requirements of Muslim computer users, this distribution also demonstrates the tremendous range of available software for Linux and how easy it is to customize Ubuntu to address the requirements of a specific audience. This distribution is currently available only for the 32-bit x86 systems. The Islamic software packages used by Ubuntu Muslim Edition can also be installed on any existing Ubuntu distribution by downloading and executing a script that installs them on your current system. For more information about Ubuntu Muslim Edition or to download a copy, see its web site at `http://ubuntume.com`. A Live CD version of Ubuntu Muslim Edition is available if you want to experiment with it before permanently installing an Ubuntu distribution.
- **Ubuntu Studio: A** version of Ubuntu that targets creating and editing multimedia. Ubuntu Studio is a relatively young Ubuntu-based distribution that was first available in late 2007. If you

routinely work with graphics, audio, or other types of multimedia, Ubuntu Studio is well worth a look. It includes standard graphics applications such as GIMP and Inkscape, but also includes Ardour 2, an audio recording and editing application much like traditional audio applications such as ProTools and Blender, for modeling and rendering three-dimensional graphics and animations; PiTiVi, Kino, and Cinepaint, for video creation and editing; and much more. I have the graphics skills of a block of wood and, music-wise, used to play the drums, but anyone with artistic ability who is interested in using open source software should definitely investigate Ubuntu Studio. You can get more information about Ubuntu Studio or download a copy from the Ubuntu Studio web site at `http://ubuntustudio.org`.

- **Xubuntu: A**n official edition of Ubuntu edition that uses the Xfce desktop rather than GNOME or KDE. Xfce is a much lighter-weight desktop environment than GNOME or KDE without sparing power or usability. Using Xfce makes Xubuntu a good choice for use on older or less-powerful systems where you still want to use a desktop environment rather than simply a window manager such as that provided by Fluxbuntu. (Desktop environments support capabilities such as drag and drop and other forms of application interaction rather than simply managing applications on the desktop — for more information about the differences between window managers and desktop environments, see the "What's a Desktop? Graphical Environments for Linux" section in Chapter 5. Xubuntu includes applications such as Thundar (a lightweight file manager), the GIMP image editor, the AbiWord word processor, the Firefox web browser, and many popular GNOME applications. For more information about Xubuntu or to download a copy, see the Xubuntu web site at `www.xubuntu.org`. A Live CD version of Xubuntu is available if you want to experiment with it before permanently installing an Ubuntu distribution.

NOTE **Canonical's Launchpad web site (`http://launchpad.net`) provides a web-based software development, source code control, and build system that hosts many other Ubuntu-based distributions, such as Elbuntu (Ubuntu with the Enlightenment window manager as its primary GUI, `http://launchpad.net/elbuntu`), Mint (focused on producing an elegant desktop, `http://launchpad.net/linuxmint`), and many more.**

Other Ubuntu-based distributions are available, but weren't listed here because I haven't actually had the chance to experiment with them or because they don't seem to be supported any longer. I did not list complete distributions such as Freespire (the community version of Linspire, `www.freespire.org/`), Linspire (`www.linspire.com`), MEPIS (`www.mepis.org/`), Sabayon (`www.sabayonlinux.org`), and others that are based on the Ubuntu packages but which are independent distributions in their own right. I also didn't list "distributions" that are only scripts that add specific sets of packages to an existing Ubuntu installation. If I've missed an installable, ready-to-run, Ubuntu-based distribution that you enjoy, please let me know. I'm always happy to learn, too.

Server Users

While it is quite possible to run server applications such as mail systems, web servers, DNS servers, print servers, and file servers on a desktop Linux system, this is often the wrong thing to do in enterprise or academic environments. Many server systems are lean and mean, command-line oriented machines that devote all available cycles to the server processes that they are running, rather than "wasting" them by running a GUI (at least by default). Ubuntu offers two distributions that are targeted to different types of servers:

- **Ubuntu Server Edition:** A version of Ubuntu that brings the power, stability, and easy management of Ubuntu Linux to the server room. By default, Ubuntu Server Edition provides a basic, secure system with no open network ports, but you can install other fundamental configurations

such as a LAMP (Linux, Apache, MySQL, and Perl) system in which these key server components are preconfigured to work together, minimizing your installation and administrative efforts. The Server Edition also provides out-of-the-box support for LTSP (Linux Terminal Server Project), which supports thin clients and enables one server system to satisfy the storage and application requirements of many lightweight desktop systems. Ubuntu Server Edition does not install a graphical user interface by default, but the standard GNOME desktop (or the X Window system window manager of your choice) is available as a package that can be manually installed. For more information about the Ubuntu Server Edition or to download a copy, see `www.ubuntu.com/products/whatisubuntu/serveredition`.

- **Ubuntu Server Edition, JeOS:** A version of Ubuntu that targets servers that are intended to run virtual machines, aka virtual appliances. JeOS stands for "Just Enough Operating System," meaning that JeOS is optimized for a small disk and memory footprint and minimal resource consumption, leaving as much of your server hardware as possible available for use by virtual machines. Unfortunately (unless you're a VMware shop), JeOS is optimized for use with the VMware virtualization solutions (`www.vmware.com`), which isn't the best virtualization solution in my opinion. I suspect that JeOS will prove equally useful for more sophisticated virtualization solutions such as Xen, KVM, and others. For more information about JeOS or to download a copy, see `www.ubuntu.com/products/whatisubuntu/serveredition/jeos`. For more information about virtualization on Ubuntu, see Bonus Chapter 7, "Using Virtual Machines and Emulators."

Specialty Platform Users

Canonical has announced a distribution called the Ubuntu Mobile Internet Device (MID) Edition that is designed to bring the power, manageability, and flexibility of Ubuntu Linux to the developers of Mobile Internet Devices (MID), which are tablet-like, touch-screen devices with full network support. The millions of developers in the open source software and Ubuntu communities are an incredible resource for embedded systems development, and MID hardware is actively under development from hardware vendors such as Intel.

Originally targeted for release in October 2007, the first releases of this distribution were based on the Ubuntu 8.04 Hardy release in April of 2008. For more information about this Ubuntu-based distribution, see `www.ubuntu.com/products/mobile`. To download install images of this distribution, see `http://cdimage.ubuntu.com/mobile/releases/hardy`. This site provides a KVM (Kernel Virtual Machine) image of Ubuntu MID that you can run as a virtual machine on any Ubuntu or Kubuntu system that can run KVM. (See Bonus Chapter 7 for more information about KVM.) This site also offers MID install images for the McCaslin and Menlow MID platforms (which are, respectively, the 2007 and 2008 generations of Intel's prototype MID hardware).

Installation Requirements

As Linux distributions have moved to the 2.6 kernel and Linux is becoming more and more popular, the chances of your having hardware that is not supported by Ubuntu Linux grow less and less. This section outlines the types of systems on which Ubuntu is supported, and the general hardware requirements for a usable system.

NOTE **One of the best things about Linux is the wide range of system types on which it is supported. Most Linux distributions, including Ubuntu, will run on older systems that would probably otherwise be discarded or used as doorstops. However (and feel free to repeat this quote), software runs slower on slower hardware. When running on older systems, you will probably want to use the command-line interface or a lighter-weight graphical environment than the default GNOME (or KDE for Kubuntu) desktops. Desktops such as Xfce or window managers such as Fluxbox and IceWM are popular and powerful alternatives to GNOME and KDE. More about these in the "What's a Desktop? Graphical Environments for Linux" section in Chapter 5.**

Supported System Types

Ubuntu and Kubuntu are supported on any of the following types of systems:

- i386 or compatible processors from Intel, AMD, and so on
- 64-bit AMD or EM64T processors (which include the Athlon64, Opteron, and EM64T Xeon)

NOTE **Ubuntu was officially supported by Canonical on PowerPC (PPC) processors, such as those used on many Macintosh computer systems, through the Ubuntu 6.10 release. Similarly, Ubuntu was also officially supported by Canonical on UltraSPARC platforms from the Ubuntu 6.10 through 7.10 distributions. Subsequently, the PPC and UltraSPARC versions of Ubuntu became community-supported distributions. This means that Canonical may still produce them, but that PPC- or SPARC-specific problems will not hold up an official Ubuntu release, and versions of Ubuntu for these platforms may not be available from all Ubuntu download sites. See the Ubuntu PowerPC FAQ at** `https://wiki.ubuntu.com/PowerPCFAQ` **for more information about Ubuntu distributions for the PPC.**

Hardware Requirements

As with any computer software, you'll have a better experience if you install and use Ubuntu on the most powerful system that you have available, but Ubuntu and Kubuntu will technically still run fine (although slowly) on your dusty 25 MHz i386. However, the American national slogan is eminently true here: "More is better." Taking off my Linux evangelist hat for a moment, you shouldn't really bother trying to install and run Ubuntu or Kubuntu on a system with a processor that runs slower than 166 MHz or that has less than 96MB of memory. I use a system with exactly those characteristics for testing purposes (an old IBM ThinkPad 380XD that I just can't bear to part with), and GNOME is excruciating on that system. If you really need to run Ubuntu on such a system, see the note earlier in this section about alternate graphical environments for low-speed or low-memory systems — or simply get Xubuntu.

The minimum hardware requirements for installing Ubuntu or Kubuntu and having a reasonable user experience are as follows:

- 700 MHz or better processor
- 384MB of memory
- CD-ROM drive
- Ethernet interface
- VGA graphics interface capable of 1024 × 768 resolution
- 3GB of available disk space

NOTE **If you want to use Ubuntu or Kubuntu's compositing and advanced visual effects capabilities, you should have at least a 1.2 GHz processor and a recent ATI, Intel, or nVidia graphics card.**

If your system satisfies or exceeds these, you're good to go. You can certainly install Ubuntu or Kubuntu on slower systems or systems with less memory, but that's like putting racing slicks on a Hyundai — you're not really going to get the most out of the experience.

Time Requirements

The amount of time that it takes to install Ubuntu or Kubuntu depends on the speed of your system, how you are configuring that system, and the type of distribution that you're installing. Installing Ubuntu or Kubuntu on a laptop that already runs Microsoft Windows, Solaris, or even Mac OS X and which you want to set up as a dual-boot machine may take an hour or two. Installing any version of Ubuntu or Kubuntu on a new machine can take less than an hour. In general, you should plan on spending an hour or two installing Ubuntu or Kubuntu — I'm assuming that you're not going to complain if it takes less time than that.

Ubuntu CDs

The DVD that is included with this book is the Ubuntu Desktop DVD, which enables you to test-drive Ubuntu on an existing computer system without changing anything, and which also provides a simple, easy-to-use installer that enables you to install Ubuntu on that system permanently. Three different CDs are available for Ubuntu. These CDs and the capabilities that they provide are as follows:

- **Desktop DVD:** The DVD that is included with this book provides a bootable version of Ubuntu Linux that enables you to run and experiment with Ubuntu without changing anything on your existing computer system. This DVD, known as a "Live DVD," also includes an easy-to-use graphical installer that makes it easy for you to permanently install Ubuntu on your computer system. Finally, this DVD includes a special installer that you can use to install and use Ubuntu on a system running Microsoft Windows without re-partitioning the disk. For information about using this DVD, see Chapter 2, "Installing Ubuntu."
- The Kubuntu web site provides an equivalent to the Ubuntu Desktop DVD that supports the same capabilities (Live CD or DVD, Installer, and Kubuntu on Windows) but uses the KDE desktop environment and provides KDE-based applications.

NOTE **When running from any live CD, any work that you do, files that you create, and so on, will be lost when you reboot your computer system unless you save it to another system over the network or to removable storage such as a USB stick, removable hard drive, and so on. See the section of Chapter 2 entitled "Using Desktop CD Persistence" for information on using a USB stick or other removable media to automatically save and restore any changes that you make while running from the Ubuntu Desktop DVD, or see the sections "Accessing Your Hard Drive from the Desktop CD" and "Copying Files to Other Machines Over a Network" in Chapter 2 for information about manually saving any work that you do while running from the Ubuntu Desktop DVD.**

- **Server Install CD:** Enables you to install versions of Ubuntu Linux targeted to machines that are being used as servers. You can choose to install a generic server and add the server software of your choice, or you can install a LAMP (Linux, Apache, MySQL, Perl) server where the traditional packages required for a Linux Web server will be preinstalled. None of the versions of Ubuntu installed from this CD include a graphical user interface, although you can always add

one subsequently. For more information about obtaining this CD and installing from it, see the section "Install Options on the Server Install CD" in Chapter 3.

- **Alternate Install CD:** Enables you to install Ubuntu on systems with certain hardware characteristics, or in specialized configurations. These include creating preconfigured systems for redistribution by Original Equipment Manufacturers (OEM), upgrading existing systems without network access, and setting up automated Ubuntu installations for multiple systems. Hardware-wise, the install options on this disk enable you to install Ubuntu on systems that use Logical Volume Management (LVM), use Redundant Arrays of Inexpensive Disks (RAID), where you want to install GRUB in a location other than the Master Boot Record (MBR), or on systems with limited amounts of memory (i.e., less than 192MB of RAM). For more information about obtaining this CD and installing from it, see the section "Install Options on the Alternate Install CD" in Chapter 3.

The Desktop DVD included with this book is the one that most people use to install Ubuntu. However, depending on the type of system that you want to create, you may want to download and burn a copy of another installation CD. The Ubuntu Web site provides freely downloadable ISO images of all of the available Ubuntu CDs, for all supported platforms, at `http://us.releases.ubuntu.com/releases`. (ISO images are files that contain an image of a CD in International Standards Organization CD format, which you can download and then burn to a CD yourself.) Pick the directory associated with the latest release, select the appropriate ISO image, download it, and burn a copy — or you can request that the folks at Ubuntu ship you a set of CDs. To do this, go to the page at `https://shipit.ubuntu.com`, create an account by entering your e-mail address and a password, and request Ubuntu CDs for the current release. You can even order free CDs for multiple system types at the same time. The Ship-It site is cool for getting copies of Ubuntu to turn on your friends, but is not a good personal alternative if you're into instant gratification, because shipping and delivery can take a few weeks.

Support for Ubuntu and Kubuntu

By its nature, computer software occasionally requires that you ask questions about how to use it or ask for help with resolving specific problems. This is especially true of software such as Linux, where you are installing not only a zillion applications, but also the operating system that they depend on. The primary advantage of off-the-shelf operating systems from a single commercial source, such as Microsoft Windows and Apple's Mac OS X, is that you can presumably contact the vendor if you're having problems installing, configuring, or using it. However, in reality, just try contacting Microsoft if you're having a problem using Windows. (Let me know how that goes.) In general, books like this one provide a central resource for installation, configuration, and general "how do I . . ." questions, but there are always specific questions that I can't anticipate. So how do you get your questions answered or find help when you need it?

Ubuntu offers an impressive array of support opportunities, ranging from community resources to paid support from Ubuntu's sponsor company, Canonical, Ltd., and a number of other companies located all over the world. The next few sections highlight the various ways in which you can ask questions, get answers, request paid support, and even hire experts to help with custom Ubuntu and Kubuntu programming and support tasks.

Community Support and Information

Because Linux software depends on the community development model, getting timely, free help for problems often relies on a similar community approach. This is one of the areas in which Ubuntu truly

shines, hosting mailing lists, blogs, and interactive forums that are all excellent sources of up-to-date information about Ubuntu. Forums and mailing lists enable you to post specific questions and receive responses from other Ubuntu users who have already solved the issue that you're experiencing. These online resources also serve as excellent feeder sites for the Ubuntu project, helping the project identify issues and common problems that should be addressed in future Ubuntu releases.

Blogs

Blogs (from the term "W*eb log*") are a popular buzzword, and it often seems as though almost anyone with a keyboard and any control over their Web site has one. The Ubuntu blog, known as Planet Ubuntu (`http://planet.ubuntulinux.org/`) is a bit different because it's not a continuous stream of consciousness from a single individual. Instead, Planet Ubuntu is a place where Ubuntu developers and community members can share various musings, insights, complaints, and successes.

If you're enough of a blog or Ubuntu fan to want to subscribe to it rather than simply visiting its Web page, Planet Ubuntu is also available in the following popular RDF (Resource Description Framework) and general markup formats:

- **FOAF:** Friend of a Friend, available at `http://planet.ubuntulinux.org/foafroll.xml`.
- **OPML:** Outline Processor Markup Language, available at `http://planet.ubuntulinux.org/opml.xml`.
- **RSS:** Really Simple Syndication or Rich Site Summary, depending whom you ask. Planet Ubuntu feeds are available in RSS 1.0 (`http://planet.ubuntulinux.org/rss10.xml`) or RSS 2.0 (`http://planet.ubuntulinux.org/rss20.xml`) formats. You should use the appropriate format for your RSS reader.

Ubuntu Forums

Forums are the latest generation of what used to be known as bulletin board systems, and are an attractive alternative to mailing lists if you have the time to visit the Web site that hosts them. Ubuntu's forums are hosted at `www.ubuntuforums.org/`. This site provides a huge selection of well-organized forums that you can easily search to find specific information, where you can post questions, or where you can simply chat with or see the posts of other Ubuntu users, dipping your toe into the waters of the Ubuntu community if you're not already an active member. You don't have to be a member of the forums to read them, but you do need to be a member to post there. Registration is free and easy — just go to `www.ubuntuforums.org/register.php`, and read and accept the Ubuntu Forum rules. You can then specify the username that you want to use on the forums, enter a password, and provide your e-mail address and some minimal personal information.

The forums index page at `www.ubuntuforums.org/index.php` displays the categories into which the Ubuntu forums are organized, as follows:

- **Absolute Beginner Talk:** A forum section where anyone can ask questions about computers, Linux, UNIX, and Ubuntu. If you were ever afraid to ask a question because it might be too basic, this forum is for you!
- **Main Support Categories:** A forum section that provides a number of different forums dedicated to various support topics for any release of any official Ubuntu/Kubuntu/Xubuntu/Edubuntu distribution. Specific forums include Apple Intel Users, Apple PPC Users, Desktop Environments, Desktop Effects & Customization, General Help, Hardware & Laptops, Installation & Upgrades,

Multimedia & Video, Multimedia Production, Networking & Wireless, Sun Sparc Users, System76 Support, Dell Ubuntu Support, and x86 64-bit Users.

- **Other Community Discussions:** A forum section that provides a number of different forums dedicated to general topics, Ubuntu and community announcements, available Ubuntu projects, colocation and development teams, and much more. Specific forums include 3rd Party Projects, Accessibility Discussions, Art & Design, Community Announcements & News, Gaming & Leisure, Ubuntu Gamers Arena (an offsite link), Education & Science, Servers & Security, Tutorials & Tips, Ubuntu LoCo Team Forums, Ubuntu Testimonials & Experiences, Ubuntu Weekly Newsletter, Development & Programming, Virtualization, and WINE.
- **Forum Community Discussions:** A set of forums dedicated to general discussions and topics about the forums themselves. Specific forums include Community Cafe, Community Cafe Games, Community Market, Forum Feedback & Help, Forum Council Agenda, Other OS Talk, The Fridge Discussions, and the Ubuntu Forum Teams.

Although other distributions have similar forum sites (such as Fedora Core's `www.fedoraforum.org` site), the Ubuntu forums embrace and reflect Ubuntu's commitment to users of the current Ubuntu release and previous Ubuntu releases that are still supported, which is truly unique.

Although you can ask questions about and discuss KDE and Kubuntu issues on the standard Ubuntu forums, Kubuntu-specific forums are available at the following locations:

- `www.kubuntuforums.net`
- `http://kubuntuway.net`

In addition to the English-language forums discussed previously, Ubuntu forums are also available in many other languages, reflecting the commitment of Ubuntu and Ubuntu users to provide a truly international Linux distribution. These are not just translated, native character set versions of the English-language forums — in many cases, different native language sites host their own forums and organize those forums differently. You can find pointers to these forums on the page at `www.ubuntulinux.org/community/forums`. At the time of this writing, specialized Chinese, Dutch, Finnish, French, German, Italian, Polish, and Portuguese forum sites are hosting forums in those languages, using any associated character sets. Ubuntu is truly an international effort!

IRC

Internet Relay Chat (IRC) is a popular mechanism for interactive online discussions of just about anything. The English-language Ubuntu IRC channel is named #ubuntu, and is available through the IRC site at `irc.freenode.net`. IRC channels for other versions of Ubuntu are available at #kubuntu, #edubuntu, and #xubuntu. Non-English IRC channels are also available, including Chinese (#ubuntu-zh), Dutch (#ubuntu-nl), German (#ubuntu-de), Hebrew and Arabic (#ubuntu-il), Italian (#ubuntu-it), Portuguese (#ubuntu-pt), Russian (#ubuntu-ru), and Spanish (#ubuntu-es), at the time of this writing. For a complete list of Ubuntu-related IRC channels, see `https://help.ubuntu.com/community/InternetRelayChat`.

An IRC channel is a great, real-time mechanism for asking about current problems and getting online help to resolve them (assuming that your problem isn't related to getting online in the first place). The Pidgin instant messaging client is installed by default as a basic part of Ubuntu Linux, which also can be used for IRC communication. Similarly, the Kopete instant messaging and IRC client is installed by default as a part of Kubuntu Linux. Providing IRC support by default makes it easy to connect and take advantage of IRC as a support and community resource. For more information, see Chapter 13, "Sending and Receiving Instant Messages."

Mailing Lists

Mailing lists are a great push format, meaning that questions and posts are delivered (pushed) directly to you, unlike forums, which are generally referred to as a pull format, because you have to connect to the Web site that hosts them and locate new posts and information yourself.

There are a huge number of Ubuntu mailing lists, many of which are quite specialized, and listing them all here would simply waste paper because you have to subscribe to them online in the first place. The standard Ubuntu mailing lists, as listed at `www.ubuntulinux.org/community/lists`, include the following:

- **Ubuntu Announcement list** (ubuntu-announce) has very few e-mails (less than one a month, usually) and will keep you up-to-date on new releases of Ubuntu, and significant new developments.
- **Ubuntu Development list** (ubuntu-devel) is intended for highly technical discussions and implementation details regarding current Ubuntu development.
- **Ubuntu Security Announcement list** (ubuntu-security-announce) is a read-only mailing list to which announcements of security updates to Ubuntu releases are posted. This list is extremely useful for Ubuntu system administrators or anyone who wants to make sure that they know about the latest security-related Ubuntu package updates.
- **Ubuntu Users list** (ubuntu-users) is an extremely high-traffic mailing list for technical support discussions and to which Ubuntu users can post new feature requests and wish lists.
- **Ubuntu Women list** (ubuntu-women) is intended as a mailing list for all Ubuntu users, volunteers, developers, and others who wish to involve more women in the Ubuntu community.
- **Kubuntu Users list** (kubuntu-users) is a high-traffic list that is focused on Kubuntu and KDE-related topics.

There are many more lists, of course — these are just some of the highlights. You can find a complete, up-to-date list of available Ubuntu mailing lists at `http://lists.ubuntu.com`, which will take you to `http://lists.ubuntu.com/mailman/listinfo`. The Ubuntu mailing lists are managed using the popular MailMan mailing list management package.



For Usenet fans, the Ubuntu mailing lists are also available as Usenet news groups thanks to the folks at Gmane (`www.gmane.org`). The Ubuntu-related mailing lists available as news groups are listed at `http://news.gmane.org/index.php?prefix=gmane.linux.ubuntu`.

What's on the Fridge?

Remember how your parents would post your latest accomplishments on their refrigerator? Let's hope they've stopped now, but the Ubuntu Web site provides a software implementation of the same concept. The Fridge (`http://fridge.ubuntu.com`) provides a central location where Ubuntu users can find out what's truly new in the Ubuntu community. It features summaries of upcoming Ubuntu-related events, recent information about Ubuntu on the Web and in print media, status messages from various Ubuntu teams and projects, and newsletters such as the Ubuntu Desktop News and Ubuntu Documentation News.

Documentation

Traditional software products provide printed or online documentation to help anticipate and answer users' questions. However, as both a writer and long-time computer user, I've always appreciated and

evangelized for good documentation. It doesn't matter how good software or an operating system is if you can't figure out how to use it. Linux documentation is an interesting issue because most Linux distributions are freely downloadable. There are few Linux distributions that you can actually buy off the shelf at a computer retailer, and even these provide relatively little printed documentation. Given the speed at which Linux distributions evolve and the tremendous variety of hardware on which Linux can be installed, complete printed documentation is difficult to produce in a timely fashion and is even harder to maintain. Novell's SUSE Linux is famous for the quality and bulk of the printed documentation that accompanies its boxed products.

Ubuntu has a large and well-organized documentation team that is focused on producing quality, user-oriented documentation that is just as easy to use as Ubuntu itself. Ubuntu documentation is available at two primary locations:

- `http://help.ubuntu.com`: The source for all of the official documentation that has been developed by the Ubuntu documentation team. By default, this page displays the documentation for the current Ubuntu release, but documentation for older releases is available through tabs that you can select in the upper-right corner of the page.
- `https://wiki.ubuntu.com/UserDocumentation`: A hierarchical collection of resources in wiki format that makes it easy to find documentation on specific topics.

A good deal of information on Ubuntu Linux and other Ubuntu-based distributions such as Fluxbuntu and Mint is available at the Ubuntu Documentation Storage facility at `http://doc.gwos.org/doku.php`.

The Ubuntu documentation team produces quality documentation that is well-organized and adheres to a single style guide to provide the sort of consistency that you'd expect from an organized documentation effort. The home page for the Ubuntu Documentation Project is at `http://doc.ubuntu.com/`, where you can find pointers to both current documents and works in progress. All communications between documentation team members is done online, using IRC and mailing lists described at `https://wiki.ubuntu.com/DocteamCommunications`. If you're interested in contributing to the Ubuntu documentation effort, a list of current and planned projects is available at `https://wiki.ubuntu.com/DocteamProjects`. Documentation that is contributed by the Ubuntu community in general is available at `https://help.ubuntu.com/community`.

Commercial Support for Ubuntu Linux

As discussed earlier, it's especially important for companies that are planning on adopting an enterprise-wide Linux solution to have a source from which they can get guaranteed support. Although you can typically find answers to most of your questions and solutions to most problems by simply searching the Web, most CEOs and IT managers won't accept "I'm googling it" as a suitable status message when an entire business is offline or some of their employees are unproductive because of a software or operating system problem. In business situations, it's important to have specific resources that you can depend on to solve problems in a timely fashion.

Although the majority of this chapter has stressed the scope and usability of the Ubuntu community and related resources in terms of helping you solve problems, commercial support is also available for Ubuntu from a variety of sources, as described in the next two sections.

Paid Support from Canonical, Ltd.

Canonical, Ltd. the sponsor of the Ubuntu Linux Project, offers two levels of paid support for Ubuntu Linux, known as 9x5 support and 24x7 support, which reflect the hours during which you can expect to

get responses to your support questions. These support packages and their current pricing are described at `www.ubuntu.com/support/paid`. These support levels are available for three different types of production systems: Desktop, Server, and Thin Client/Cluster systems, each of which has different costs depending on the hours of support that you're interested in.

The software packages that Canonical supports depends on the portion of the Ubuntu software repository in which those packages are located. As described in detail in Chapter 23, "Adding, Removing, and Updating Software," the software in the Ubuntu repository is grouped into several different classes, essentially depending upon the license(s) under which a software package has been released. Ubuntu support agreements include full support for packages in the *main* class, partial support for packages in the *restricted* class, and no support for software in the *universe* and *multiverse* classes. See Chapter 23 for details on the organization of the Ubuntu repository and the differences between package classes.

The Ubuntu Marketplace

The Ubuntu Marketplace is a portion of the Ubuntu Web site that lists the network of companies that provide support for desktop and server systems running Ubuntu Linux. The main page for the Ubuntu marketplace is `www.ubuntu.com/support/supportoptions/marketplace`, which provides centralized access to lists of companies all over the world that support Ubuntu, organized into separate pages listing such companies in Africa, Asia, Europe, Latin America, North America, and Oceania. (A related URL is the Community Market forum at `www.ubuntuforums.org`, which was mentioned earlier, although that contains much more information than the actual Ubuntu Marketplace.)

Providing a centralized clearinghouse for companies that can help you or your firm with support problems is a tremendous advantage for companies that are just moving to Linux and would like to engage with a support organization that is geographically close to your physical location.

At the time of this writing, there are no specific certification requirements to have your company listed in the Ubuntu Marketplace — you simply fill out a form that is linked to from the Marketplace page.

Getting More Information About Ubuntu

Ubuntu Linux is increasing in popularity faster than any Linux distribution that I have previously encountered. Technical excellence aside, much of the credit for its increasing popularity lies in the excellent organization and breadth of coverage provided on the main Ubuntu Web site. However, let's face it — the Ubuntu Web site is largely blowing its own horn. What do other people say? Are there locations other than the mother ship where you can go for information about and help with Ubuntu?

The answer to these questions is a definite "yes!" There are a huge number of sites that provide information about Ubuntu beyond simple software reviews. In addition to the support and general Web resources listed previously in this chapter, some of my favorite Ubuntu-related sites are as follows:

- **DistroWatch** (`http://distrowatch.com/table.php?distribution=ubuntu`) provides summary information about the contents of most Linux distributions, including Ubuntu. Their Ubuntu page provides high-level information about the contents of the various Ubuntu releases, but more importantly provides links to many Ubuntu-related Web sites, reviews of the various Ubuntu releases, and much more. If you're not already familiar with DistroWatch, it's an essential Linux site for finding out almost anything about any Linux distribution.
- **Ubuntu Blog** (`http://ubuntu.wordpress.com/`) is a blog about Ubuntu that provides a great selection of entries about general Ubuntu tasks organized into categories such as administration,

office, servers, and so on. It also features links to other sites and great task-specific articles such as the best instructions for getting the MythTV package working on an Ubuntu system.

- **UbuntuGuide** (`http://ubuntuguide.org/`) hosts an Ubuntu "Getting Started" guide that is an excellent information resource, although it may not always reflect the latest Ubuntu release. The Ubuntu Starter Guide provides a great deal of very detailed information about how to do specific tasks on Ubuntu Linux, and is well worth a look.
- **Ubuntu Women** (`www.ubuntu-women.org/`) provides FAQs, a wiki, a blog, and mentoring programs, and is much more focused on getting more women involved in Ubuntu and FLOSS (Free/Libre/Open-Source Software) in general. This is a great site with a great message and purpose, through which we all win.
- **Ubuntux** (`www.ubuntux.org/`) is a community of Ubuntu users that provides a variety of forums, blogs, links to recent articles about Debian and Ubuntu, and a variety of other resources designed to help users work with Ubuntu. The forums are especially nice, covering the spectrum of topics from getting started with Ubuntu to specific customization and optimization topics. Ubuntux also features forums targeted toward Ubuntu-based distributions such as Kubuntu, Edubuntu, and Ubuntu Lite, as well as a forum on the Ubuntu server distribution. The Ubuntux site also offers an RSS feed to help you keep up-to-date with the latest Ubuntu-related happenings.

There are many other Ubuntu-related sites on the Web, with more appearing every day. The DistroWatch site provides a good collection of Links to Ubuntu-related sites and reviews, but as with anything on the Internet, your favorite search engine is your friend and will quickly help you find hundreds of other sites to search for answers to specific questions or simply to see what others think and say about Ubuntu.

Summary

Ubuntu is the fastest growing Linux distribution in recent memory, and is one of the finest examples of the power of open source and community that I've ever seen. After providing some general information about Linux, this chapter provided an overview of the philosophy behind Ubuntu Linux and the goals of the distribution. As discussed in this chapter, much of the success of Ubuntu to date beyond its technical excellence and ease of use is rooted in a rich, fast-growing user community and a well-organized Web site that provides easy access to various Ubuntu-related resources.

Chapter 2

Installing Ubuntu and Kubuntu

This chapter discusses how to install or simply experiment with Ubuntu or Kubuntu Linux on a new or existing 32- or 64-bit PC system. This book includes a DVD for the latest version of Ubuntu Linux at the time of this writing — sorry, Kubuntu fans! As I explain in Chapters 7 ("Using the KDE 3 Desktop") and 9 ("Managing E-mail and Personal Information with Evolution"), Kubuntu is somewhat in flux at the moment, and it was therefore impossible to pick just one Kubuntu release to include. If you want to use Kubuntu, I'll explain how to obtain an equivalent CD for the version of Kubuntu that's right for you. The Ubuntu and Kubuntu Desktop discs both work in exactly the same way, except that they feature different desktop environments and slightly different installers, so all of the information in this chapter is relevant regardless of whether you're installing or test-driving Ubuntu or Kubuntu.

The DVD included with this book, the Ubuntu 8.10 Desktop DVD, is an example of a Live CD, which is a bootable CD (or DVD, in the case of the disc included with this book) that provides a great way to experiment with a preconfigured Ubuntu installation without actually modifying your hard drive or going through the installation process. You can use Live CDs as a portable Ubuntu or Kubuntu installation that you can take with you, temporarily converting other computer systems to Ubuntu or Kubuntu Linux systems regardless of whether the machine is currently running Microsoft Windows or another Linux distribution. The Live CDs also offer an impressive Live CD Persistence function that enables you to save work and system configuration done while using the Desktop CD to removable media that you can then reuse the next time you boot from the Live CD. In most cases, you can also mount and access your existing disk partitions after booting from the Ubuntu or Kubuntu Desktop CD, so that you can even do real work while using the Live CDs and save it to the partitions used by your other operating system(s).

The Ubuntu and Kubuntu Desktop CDs take the traditional Live CD concept one step further by also enabling you to quickly install Ubuntu or Kubuntu Linux using an easy-to-use installer that is available as an icon on the desktop that you see after booting from the Live CD. The Ubuntu installer can

IN THIS CHAPTER

Booting a Desktop CD

Installing Ubuntu or Kubuntu

Test-driving Ubuntu or Kubuntu

Accessing existing partitions from a Desktop CD

Creating bootable USB drives

Installing Ubuntu or Kubuntu without partitioning

automatically reduce the size of existing Microsoft Windows or Linux partitions to free up space in which you can install Ubuntu (assuming that your existing partitions aren't full, of course).

But wait, there's more! Inserting the Live CD into a system that is running Microsoft Windows enables you to use a package known as Wubi (Windows-based Ubuntu Installer) to install Ubuntu or Kubuntu on your Windows system without modifying that system in any way, except for adding an option to boot Ubuntu or Kubuntu when you turn it on (and, of course, using some of your existing disk space to store the Ubuntu or Kubuntu boot image). More about that later!

If you're using a 32-bit x86 system and want to run or experiment with Ubuntu Linux, the Ubuntu Desktop DVD provided with this book is for you. If not, this chapter explains how to get the Ubuntu or Kubuntu Live CD that's appropriate for the system you're using. This chapter explains how to install Ubuntu or Kubuntu Linux from a Desktop CD or test-drive Ubuntu or Kubuntu from a Live CD without changing your existing system. If you're simply experimenting with Ubuntu or Kubuntu from a Live CD, this chapter explains how to access your existing disk partitions and files so that you can work with them, and explains how to preserve any work that you've done while using the Desktop CD. It concludes by discussing how to use Wubi to install Ubuntu or Kubuntu on an existing Microsoft Windows system.

NOTE **Versions of Ubuntu are also available for the PowerPC (PPC) and UltraSPARC platforms, but these are no longer officially supported by Canonical. Ubuntu for PPC is a community-supported project that is still quite active — see** `http://ubuntuforums.org/showthread.php?t=427714` **for information about downloading and installing Ubuntu on a PPC system. Ubuntu for UltraSPARC systems seems to be somewhat less active — see** `http://ubuntuforums.org/forumdisplay.php?f=146` **for information about finding and installing Ubuntu on your favorite doorstop, er, I mean SPARC system.**

Ubuntu, Kubuntu, and Which Kubuntu

The section of Chapter 1 entitled "Which Ubuntu Should I Choose?" provided an overview of Ubuntu-based distributions, including Kubuntu. This book focuses on Ubuntu and Kubuntu, but you still have to choose one or the other (and if you choose Kubuntu, you have to choose between the KDE 3-based Kubuntu 8.04 or the KDE 4-based Kubuntu 8.10 and later). This section provides some guidance for making your initial decision. My suggested decision tree for figuring out which Ubuntu or Kubuntu version to install is as follows:

- If you really don't care, skip to the next section and install Ubuntu from the DVD that accompanies this book, a 64-bit download if you have a 64-bit system, or whatever the latest version of Ubuntu is by the time that you read this. The standard Ubuntu distribution is my personal favorite. You'll still be able to install and use the KDE applications discussed in this book on a system running the standard Ubuntu distribution.
- If you are familiar with Linux and have a strong preference for a specific desktop environment, install Ubuntu if you want to run GNOME as your desktop environment, or install Kubuntu if you want to run KDE as your desktop environment. If you want to run KDE, you have a couple of options:
 - If you want to stick with the mature KDE 3 desktop, you'll want to install Kubuntu 8.04. However (and unfortunately), let's face facts — KDE 3 is an evolutionary dead end. It is a great desktop, is quite stable, and is also no longer actively under development. Everything that works now will always work (and you'll get updates for the supported lifetime of Kubuntu 8.04), but you won't see many new KDE 3 applications.

- If you want to run the version of KDE that is the future of KDE, you'll want to install Kubuntu 8.10 (or later). KDE 4 is still on the computer equivalent of a shakeout cruise, and some things may not work as well or as easily as you'd like them to. However, they will be fixed, and you'll always be able to use the latest and greatest KDE applications if you install Kubuntu 8.10 or later.

- If you are new to Linux but expect to depend on friends for support or assistance, ask them whether they are running Kubuntu or Ubuntu. (If they're running some other Linux distribution, encourage them to upgrade to Ubuntu or Kubuntu. This book makes an excellent gift for almost any occasion.) If they are using a GNOME-based system, you'll probably want to install a standard Ubuntu distribution such as the one that is provided on the DVD that accompanies this book. If they are using a KDE 3-based distribution, you will probably want to install Kubuntu 8.04. If they are using a KDE 4-based distribution, you will probably want to install Kubuntu 8.10 (or whatever the latest version is). You can still install and use any of the GNOME applications that are discussed in this book on a system running either version of Kubuntu.
- If you are new to Linux and are flying solo or could care less about the alphabet soup of Linux desktops, install Ubuntu. After all, an Ubuntu DVD came with this book — if you are installing on a 32-bit system, all you need to do to get started is to pop it in your computer's DVD drive, and away you go. You'll still be able to install and use the KDE applications discussed in this book on a system running the standard Ubuntu distribution.

Don't view this choice as an irritating decision, like which brand of coffeemaker or type of tires to buy. View this choice as a kid-in-a-candy-shop decision, where you simply have to decide which delicious treat to eat (or a Jaguar/Mercedes/BMW decision for the more mature or diabetic). One great aspect of Linux distributions is that they're all free, so there is absolutely no cost to you (except your time) if you try Ubuntu, decide you don't like GNOME, and then want to try Kubuntu 8.04 or later.

Getting a 64-Bit or Kubuntu Desktop CD

The Ubuntu Desktop DVD provided with this book will boot on any x86 PC system. If you want to install or experiment with Kubuntu, or want to install Ubuntu on a 64-bit computer system, you will need to obtain a different, but equivalent, Desktop CD. If you have a reasonably fast Internet connection (or plenty of time) and a CD burner, the easiest way to get a different Desktop CD is to download an ISO image. (An ISO image is a file containing an image of a CD-ROM or DVD that is in the format mandated by the International Standards Organization 9660 specification.) To download the right ISO image for you, do one of the following:

- To get a 64-bit Ubuntu Desktop CD, go to Ubuntu's Web page at `www.ubuntu.com/download`, select the version of Ubuntu that you want a CD for, select a mirror site in your country or one that's closest to you in general, and select an ISO image of the 64-bit Desktop CD.
- To get a 32- or 64-bit Kubuntu Desktop CD, go to the Kubuntu Web page at `www.kubuntu.org/getkubuntu`, select the version of Kubuntu that you want a CD for, select a mirror site in your country or one that's closest to you in general, and select the ISO image of the Desktop CD that is appropriate for your computer system.

Save the downloaded ISO image to a file, burn that ISO image to a CD once the download completes, and you're ready to go!

If you have a slow Internet connection or do not have a CD burner, don't panic! As mentioned in the introduction for this book, the Ubuntu folks will even ship you Ubuntu CDs, although they take a few weeks because they're shipped from the Netherlands. This isn't the right solution for those of us who are into instant gratification, but it might work for you. To request CDs direct from Ubuntu, go to the Ubuntu Web page at `https://shipit.ubuntu.com`, create an account by entering your e-mail address and a password, and request one or more sets of Ubuntu CDs for the current release. You can even order free CDs for multiple system types at the same time.

Unfortunately, the Kubuntu folks don't have the resources to be able to package and ship CDs at the moment, so if you have a pitifully slow Internet connection or don't have a CD burner, you'll have to prevail on one of your friends to download and burn a Kubuntu CD for you. A blank CD and a pizza can be quite persuasive in this situation.

Booting from a Desktop CD

Booting from an Ubuntu or Kubuntu Desktop CD is as easy as inserting the appropriate Desktop CD into your system's CD drive, restarting your system, and making sure that the system boots from the CD drive before booting from a hard disk partition.

CAUTION **If your PC boots from the hard drive even if the correct Desktop CD is present in your CD drive, you'll need to modify the *boot order* in your BIOS settings. The boot order is the sequence in which available devices are searched for bootable discs when your system is powered on. To do this, turn your computer on and press the key on your keyboard that gives you access to the BIOS. This is typically either the Del (Delete) or F2 key on most modern systems — if not, the key that you'll need to press is usually identified at the bottom left of your screen when you turn on your computer and it first starts up.**

Depending on the type of BIOS your computer uses, boot order settings are usually stored in an Advanced Settings or Boot screen, which you can navigate to using the arrow keys on your keyboard. Press Return to display this screen once its name is highlighted. Once this screen displays, use the down-arrow key to navigate to the First Boot Device or CD Drive entry, and see the help messages at the right side of the screen for information about how to make your CD Drive the first boot device. You can then press the Escape key to exit this screen, and press F10 to save the new settings, exit the BIOS settings screen, and reboot.

When you boot from the Ubuntu Desktop CD, you'll initially see a screen like the one shown in Figure 2-1.

This screen enables you to specify the language that you want to use when booting or installing Ubuntu or Kubuntu from the Desktop CD. This highlights one of the great features of Ubuntu and Kubuntu — they truly are designed for and committed to an international audience.

To continue, use the cursor keys to select your language and press Return. If you do not select a language within 30 seconds, English will be used as the default language. The main boot screen for the Live CD displays. Figure 2-2 shows the main boot screen for the Ubuntu Live CD.

This screen enables you to select whether you want to boot from the Desktop CD to try Ubuntu, install Ubuntu, test the CD for defects, test your system's memory, or ignore the Desktop CD and boot from the system's first hard drive.

This screen is essentially identical if you are installing Kubuntu 8.04 or 8.10, providing exactly the same options but displaying the Kubuntu logo and using Kubuntu blue rather than Ubuntu brown. Figure 2-3 shows a Kubuntu Desktop CD's main boot screen.

FIGURE 2-1

Specifying a language when booting the Desktop CD



FIGURE 2-2

Boot options from the Ubuntu Desktop CD

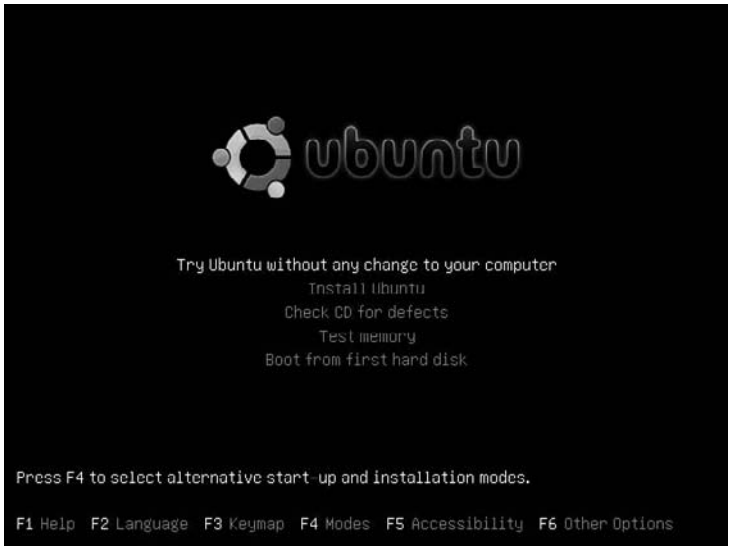


FIGURE 2-3

Boot options from a Kubuntu Desktop CD

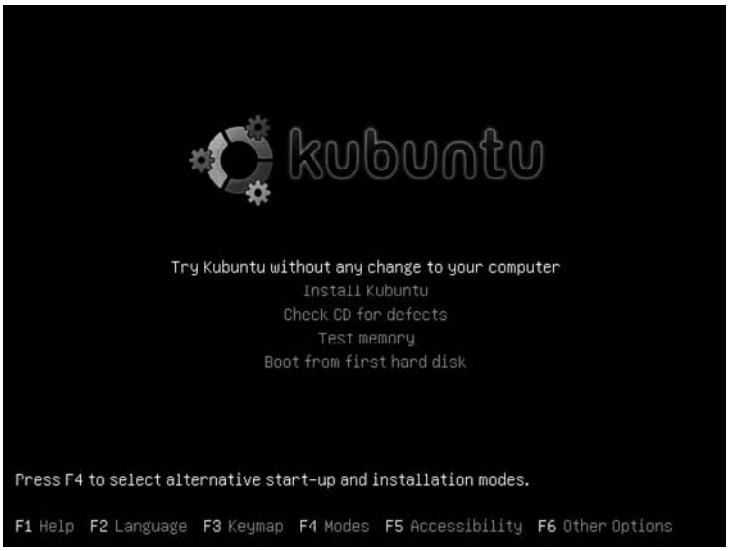


FIGURE 2-4

The standard Ubuntu desktop from the Desktop CD



To boot Ubuntu or Kubuntu Linux on your system, press Return or wait 30 seconds, at which point your system will automatically boot from the Desktop CD. If you're booting from an Ubuntu Desktop CD, the standard Ubuntu Desktop displays, as shown in Figure 2-4.

If you're booting from a Kubuntu 8.04 Desktop CD, the standard KDE 3-based Kubuntu desktop displays, as shown in Figure 2-5.

If you're booting from a Kubuntu 8.10 (or later) Desktop CD, the standard KDE 4-based Kubuntu desktop displays, as shown in Figure 2-6.

Once you're running Ubuntu or Kubuntu from the Desktop CD, you can:

- Install the distribution from the Desktop CD as described in the next section.
- Explore Ubuntu using the menus that are available in the upper-left corner of the screen and the sample files provided in the Examples folder.
- Explore Kubuntu using the menus that are available in the lower-right corner of the screen.

Experimenting with Linux after booting from a Desktop CD is described in the section "Test-Driving Ubuntu or Kubuntu Linux" later in this chapter. When test-driving Ubuntu or Kubuntu, you can practice running specific applications and even experiment with installing and configuring the software that is discussed in the remainder of this book.

FIGURE 2-5

The KDE 3-based Kubuntu desktop from a Kubuntu 8.04 Desktop CD



FIGURE 2-6

The KDE 4-based Kubuntu desktop from a Kubuntu 8.10 Desktop CD



Installing Ubuntu or Kubuntu from a Desktop CD

There are two ways to use a Desktop CD to install Ubuntu or Kubuntu on a computer system:

- Boot from the CD, set your language, use the arrow keys to select the Install Ubuntu option, and press Return or Enter on your keyboard.
- Boot from the CD, set your language, and press Return or Enter on your keyboard to boot using the "Try Ubuntu without any change to your computer" option. Once your system boots using the Desktop CD, double-click the Install icon on the desktop.

These two mechanisms run the same installer to permanently install Ubuntu or Kubuntu on your computer system (depending on the Desktop CD that you booted from). The instructions in this section describe how to install the Desktop, or end-user, version of Ubuntu or Kubuntu from an Ubuntu or Kubuntu Desktop disc, such as the Ubuntu Desktop DVD that is included with this book. See Chapter 3, "Installing Ubuntu on Special-Purpose Systems," for information about installing other versions of Ubuntu Linux, such as the Ubuntu Server or OEM distributions, which require that you obtain a different installation CD from the Ubuntu Web site. For an overview of the different Ubuntu CDs, see the section of Chapter 1 entitled "Ubuntu CDs."

The Ubuntu and Kubuntu installers are very similar, displaying equivalent screens and prompting you for equivalent information during the installation process. The screenshots in the rest of this chapter show the

Ubuntu installation process except where the installers differ to any significant extent. Don't worry — if the screens or procedures differ in any meaningful way, I'll show both.

CAUTION **If you are installing Ubuntu or Kubuntu onto a system on which another operating system is already installed, make sure that you have backed up any data that you want to preserve to removable media such as a CDROM, DVD, USB flash drive, and so on. If you want to continue using that operating system, the Ubuntu and Kubuntu installers do a great job of automatically repartitioning your disk so that you can boot your computer system into either your existing operating system or Ubuntu/Kubuntu, but nothing is perfect, and "soft" is a big part of "software." It is completely impossible to have too many backups of your critical data. If your computer system already has another operating system installed on its hard drive(s) that you want to preserve, the Ubuntu and Kubuntu installers will automatically offer to resize the existing operating system partitions and install Ubuntu or Kubuntu as an alternate operating system on your machine.**

You'll need to supply a few pieces of information during the installation process so that Ubuntu or Kubuntu is correctly configured for your system, language, and geographic location. The installation process is as follows:

1. After starting the installer using either of the mechanisms described at the beginning of this section, the dialog shown in Figure 2-7 displays, which prompts you to select the language you would like the installation process to use in subsequent dialogs and messages. You can either click Forward or press Return to accept the default value of English and proceed to the next screen, or use the arrow keys to scroll up and down in the list of supported languages until you find your native language, and then click Forward or press Return to proceed to the next screen.

FIGURE 2-7

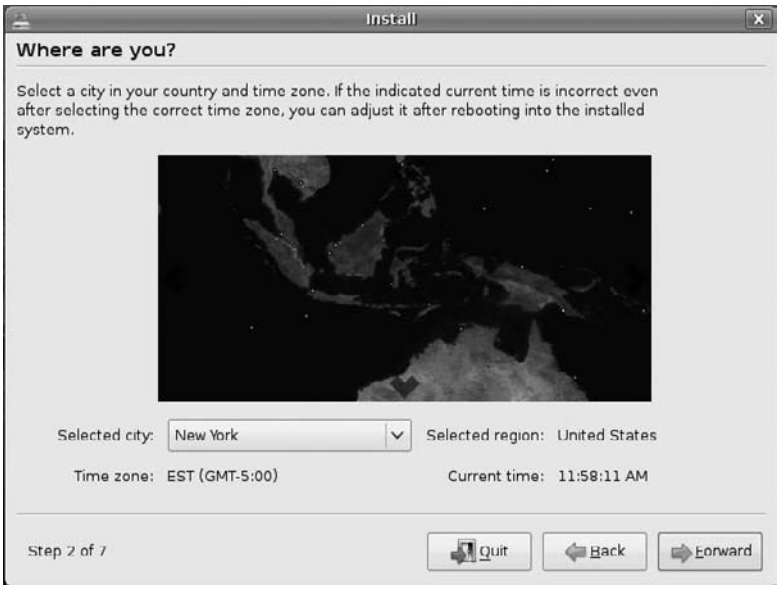
Specifying the language for the installation process



2. The dialog shown in Figure 2-8 prompts you to specify the location where you are using your computer. This value is used to further customize the language that you specified in the previous screen by using the appropriate date and time formats, numeric conventions, and currency values for your geographic location. If the city shown as the Selected City is incorrect, click the name of that city to display a drop-down menu that enables you to select a major city that is close to your geographic location and therefore is in the same time zone as you are. You can also simply use the mouse to select a nearby city from the map displayed in the center of this dialog. Once a correct city is set, click Forward or press Return to proceed to the next screen.

FIGURE 2-8

Specifying your geographic location



3. The dialog shown in Figure 2-9 prompts you to specify the type of keyboard attached to your computer. The default value is based on your settings on previous screens — if this is incorrect, select your keyboard type from the full list of supported keyboard types. You can enter sample text into the text area at the bottom of the screen to verify that the correct symbols are being displayed as you type. Once you are sure that the correct keyboard is set, click Forward to proceed.
4. The dialog shown in Figure 2-10 enables you to select where you want to install Ubuntu or Kubuntu. This dialog looks different and displays different options depending on how the disks in your system are currently configured and what partitions they already contain (if any). Figure 2-10 shows the screen that you would see if you were installing Ubuntu on a system where no operating system is currently installed.

 This dialog also looks different when installing Ubuntu and Kubuntu. Figure 2-11 shows the same dialog when installing Kubuntu 8.10.

FIGURE 2-9

Specifying your keyboard type

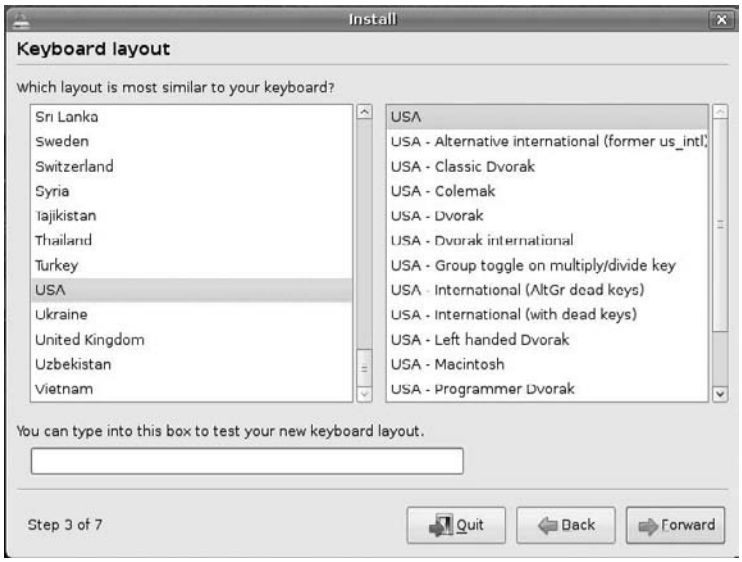


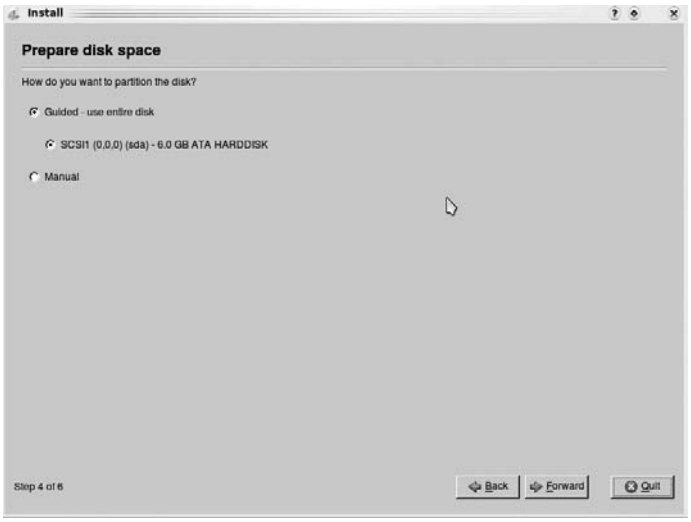
FIGURE 2-10

Specifying where you want to install Ubuntu Linux



FIGURE 2-11

Specifying where you want to install Kubuntu Linux



The option that you want to select on this screen therefore depends on the current state of your system and where you want to install Ubuntu or Kubuntu. If your system contains multiple hard drives, they will all be listed in this dialog. Your options are as follows:

- If you are installing on a new computer system or hard drive, the "Guided — use entire disk" option will be selected, as shown in Figures 2-10 and 2-11.
- If you are installing Ubuntu on an x86 or 64-bit PC system on which an operating system is already installed and do not have unallocated space on your disk, the "Guided — resize . . ." option will be selected. Both the Ubuntu and Kubuntu installers will offer to resize your current partition and install Ubuntu or Kubuntu into the space that is freed. The Ubuntu installer's version of this screen is shown in Figure 2-12.
- If your system already contains an operating system but you do not want to preserve it, you can select the "Guided — use entire disk" option to erase your existing operating system and install Ubuntu or Kubuntu on the entire disk on which it was located. The suggested resizing will still be displayed, but will be de-selected, as shown in Figure 2-13.
- If your system contains multiple hard drives, they will all be listed in this dialog. Select the disk on which you want to install Ubuntu or Kubuntu.
- If you have unallocated space on your disk (such as if you follow the instructions about manually setting up a dual-boot system later in this chapter), select the "Guided — use the largest contiguous free space" option.
- If you are an experienced Linux user and want to manually specify how to partition your disk, select the Manual option and see the sidebar entitled "Manually Partitioning a Disk" for additional information about using the manual partitioning tools provided by the Ubuntu and Kubuntu installers.

FIGURE 2-12

Suggested partition resizing in order to install Ubuntu

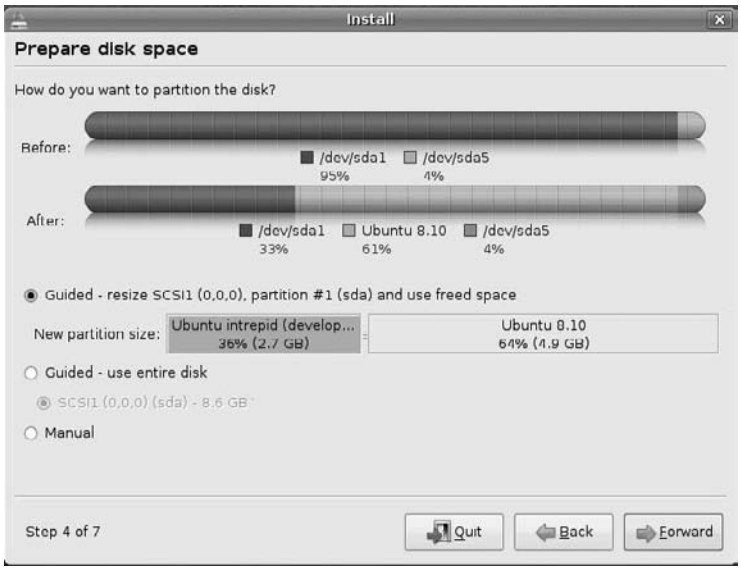
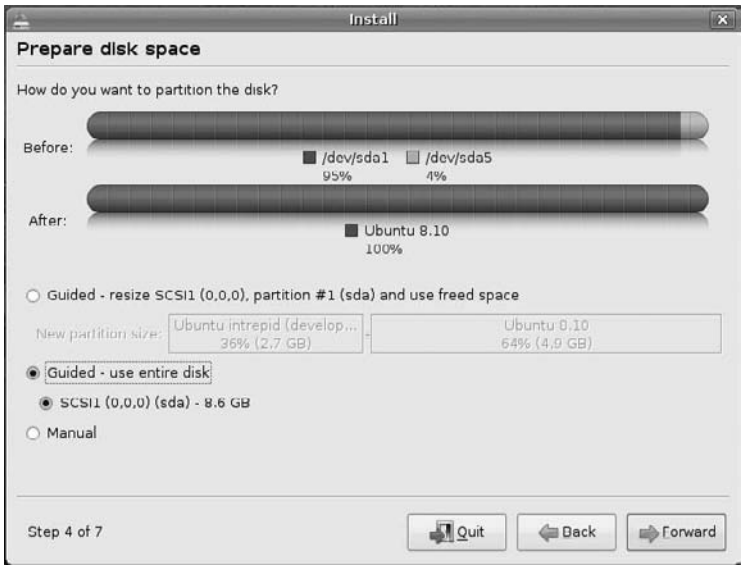


FIGURE 2-13

Deleting an existing operating system in order to install Ubuntu



Once you have selected the appropriate option for your system, click Forward to continue.

5. The dialog shown in Figure 2-14 enables you to provide personal information that will be used to create your account during the installation process. This account will have the ability to execute privileged commands on the system, as explained in the section of Chapter 4 entitled "Performing Privileged Operations in Ubuntu." After you have filled out all of the fields on this form, click Forward to proceed.

FIGURE 2-14

Specifying information about your account



6. The screen shown in Figure 2-15 displays if another Linux operating system was detected on your system and you are preserving that operating system. You will not see this screen if you are installing Ubuntu or Kubuntu on a new system or are installing on a system where you are deleting any previously installed operating systems.

 If any customized account information, such as e-mail settings, browser bookmarks, and so on, were found in that operating system installation, this dialog offers to copy them to the new installation that you are creating.

 If no custom settings were found, the dialog simply reports this, and you can click Forward to proceed to the final installer screen. If any account information was detected and you want to copy that information, select the account(s) in the dialog shown in Figure 2-15 and click Forward to proceed to the final installer screen.

7. The Ubuntu desktop installer displays the screen shown in Figure 2-16, which summarizes the information that you have supplied on previous dialogs. Click Install to proceed.

FIGURE 2-15

Importing existing account information



FIGURE 2-16

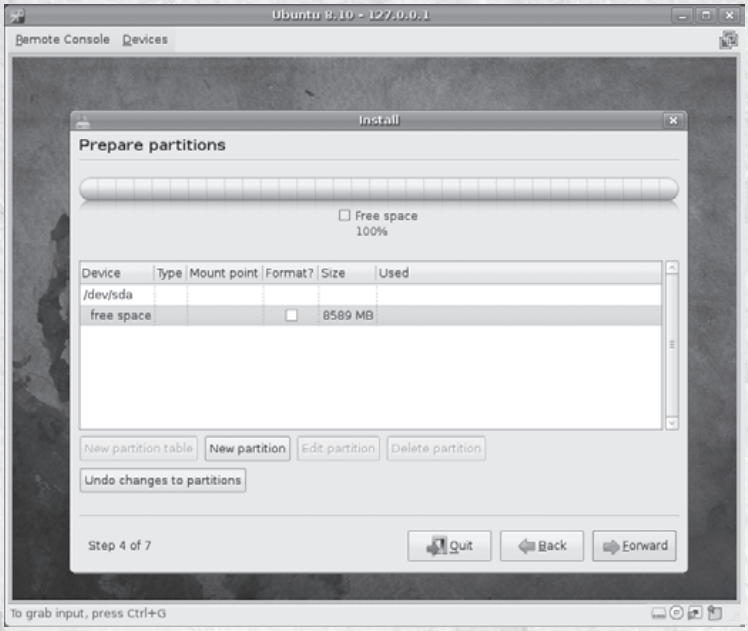
The summary dialog for starting the installation process



Manually Partitioning a Disk

If you are an experienced Linux user and have a specific reason to do so, you may want to manually specify how you want to divide up the space on your hard disk (or unallocated space on a dual-boot system). If you do not manually partition your disks, Ubuntu and Kubuntu provide reasonable default values for the different partitions that it creates automatically (a swap partition whose size is a multiple of the amount of memory in your system and a single large root partition). However, you may want to further divide the disk space on your system and/or allocate it differently. For general information about disks and partitions, see the section of Chapter 4 entitled "Disks, Partitions, and Mount Points" for more information.

If you are installing Ubuntu on a new disk and select the Manual option on the dialogs shown in Figure 2-10, 2-12, or 2-13, a dialog like the one shown in the following figure displays after you click Forward.

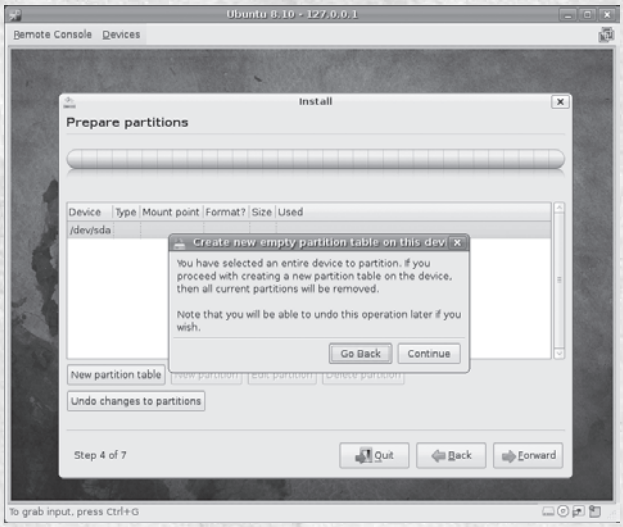


Ubuntu's manual disk partitioner

If any existing partitions are present on the selected disk, they will be shown in this dialog. If you are installing Ubuntu on a new disk, you will be prompted to create a new partition table on the disk, as shown in the following figure.

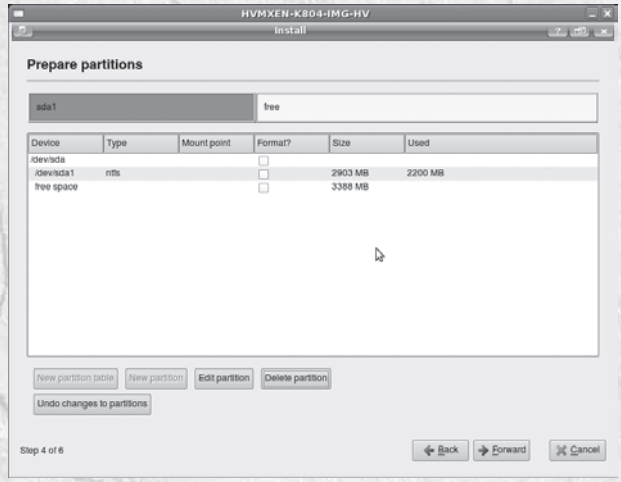
continued

continued



Creating a partition table on a new disk when installing Ubuntu

If you are installing Kubuntu and select the Manual option on the dialog shown in Figure 2-12, you will see a dialog like the following:



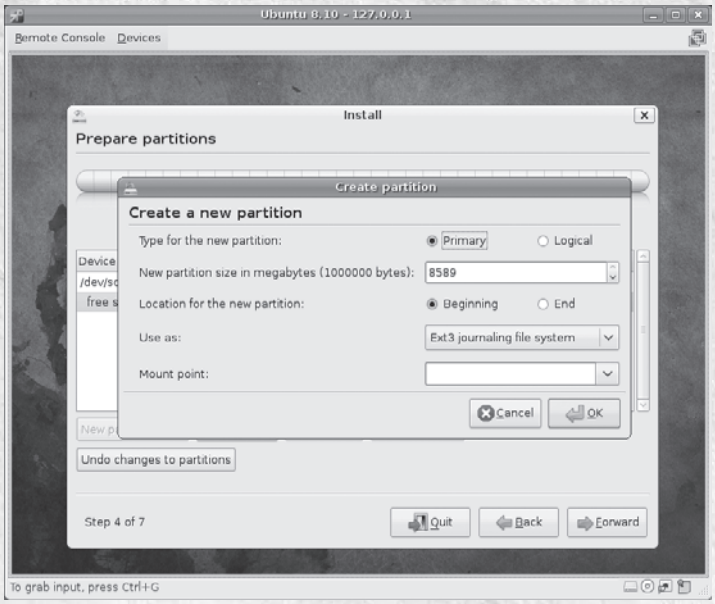
Kubuntu's manual disk partitioner

continued

continued

This screenshot was taken on a different system, and therefore shows existing partitions. If you are installing Kubuntu on a new disk, you must manually create a new partition table for the disk by clicking the New partition table button before proceeding.

If your disk contains existing partitions and you want to delete any or all of them on your disk, select each partition that you want to delete and click Delete partition to remove it and return the space associated with that partition to the pool of free space that is available on your disk. Be careful — deleting a partition permanently removes any information that it contains. Next, select the free space on your disk and click New to define a new partition. If you are installing Ubuntu, the dialog shown in the following figure displays:



Specifying partition size during manual partitioning on Ubuntu

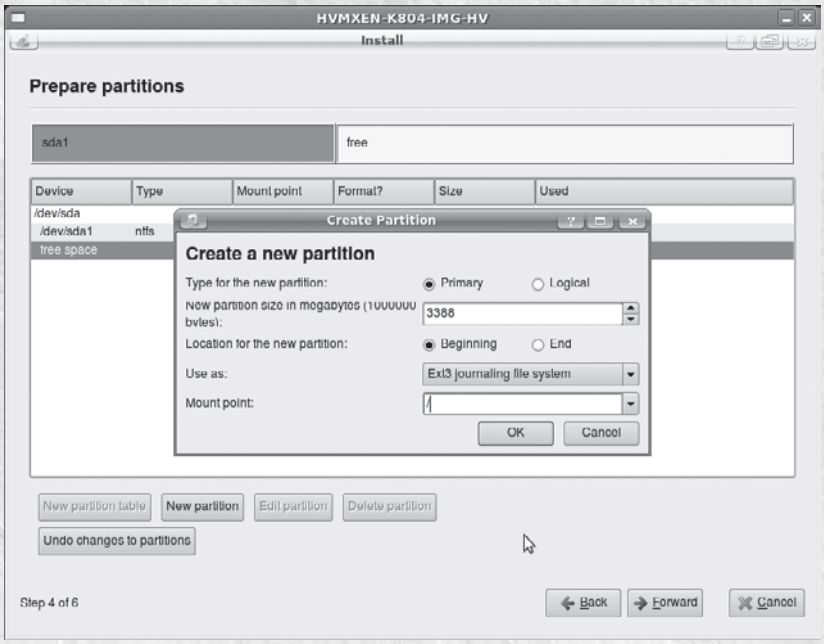
If you are installing Kubuntu, an almost identical dialog displays, as shown in the following figure:

Enter the size of the partition that you want to create and verify the type of partition that you are creating (physical or extended), along with its format (usually ext3 or swap) and the location at which you want to mount that partition; click OK. Depending on whether you are installing Ubuntu or Kubuntu, this redisplays the dialog shown in the first or second figure in this sidebar, which now contains information about the partition that you just defined.

Repeat this step for every partition that you want to create. You must create at least one swap partition, which Ubuntu or Kubuntu will use to support virtual memory and which should be a multiple of the amount of memory in your system, and one partition to hold your system's root filesystem, which should be at least two gigabytes in size. Once you have defined all of the partitions that you want to create, and you're sure that these settings are correct, click Forward to proceed.

continued

continued



Specifying partition size during manual partitioning on Kubuntu

8. The installer displays several progress dialogs as it formats the partitions on your hard disk, installs Ubuntu or Kubuntu Linux to your system, creates accounts and other system configuration information, and so on. Once the installation completes, a dialog like the one shown in Figure 2-17 displays.

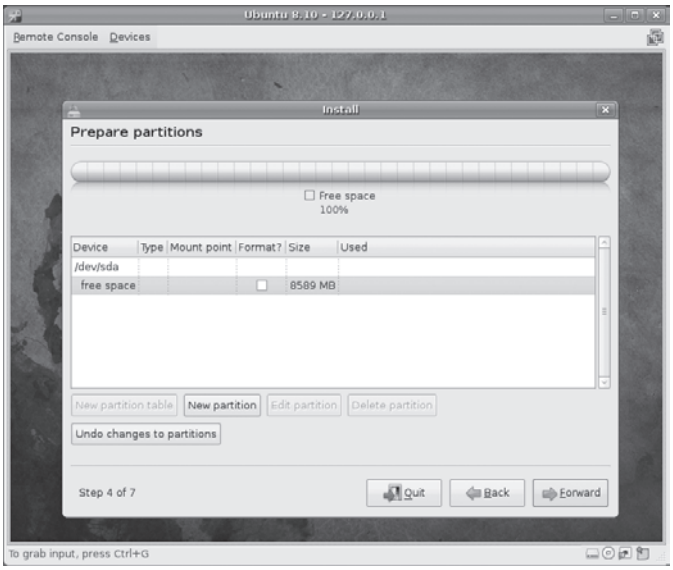
This completes the Ubuntu or Kubuntu Desktop installation process — the next time that you boot your system (after removing the Desktop CD), you're ready to run Ubuntu or Kubuntu Linux from your system's hard drive(s). Congratulations!

Booting Ubuntu or Kubuntu

If you installed Ubuntu or Kubuntu as the only operating system on a new or existing computer, you will see a quick message about "GRUB loading, please wait . . . Press ESC to access the menu" on the screen once the system begins to boot from the hard drive. Ubuntu or Kubuntu will boot automatically after 5 seconds, but you can press the Escape key during this time to display the GRUB boot menu, which enables you to boot different kernels (if more than one kernel is installed), boot any available kernel in different modes, or test system memory. For more information about GRUB and its menus, see the section of Chapter 22 entitled "The Boot Loader."

FIGURE 2-17

The final installation dialog



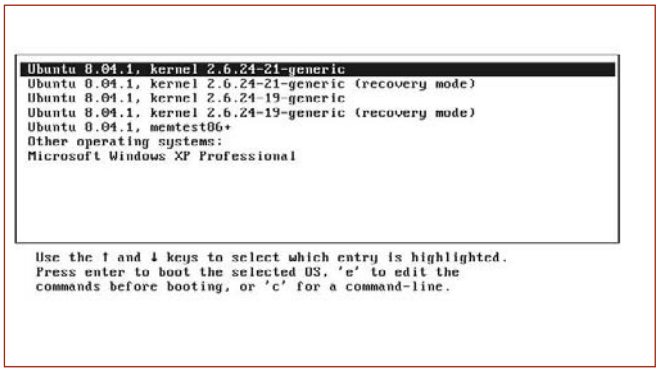
Ubuntu's manual disk partitioner

Booting Ubuntu or Kubuntu on Dual-Boot Systems

If you installed Ubuntu or Kubuntu on a system on which another operating system is also installed, the GRUB boot loader menu displays automatically to let you select the operating system that you want to boot. By default, your dual-boot system will boot Ubuntu or Kubuntu after 30 seconds. On an x86 system, this screen looks like the one shown in Figure 2-18.

FIGURE 2-18

Selecting an operating system on a dual-boot x86 system



The First Time You Boot Ubuntu or Kubuntu

The first time that you boot Ubuntu or Kubuntu, you will want to update your system. This ensures that you have the latest and best versions of all of the standard Ubuntu or Kubuntu software, which provides security-related fixes, bug fixes, and general enhancements. The Ubuntu or Kubuntu update manager should run automatically and display an icon alerting you that updates are available in the top-right portion of the panel on an Ubuntu system, or the bottom-right portion of the Kicker or Kickoff panel on a Kubuntu system. For information about updating an Ubuntu system, see the section of Chapter 23 entitled "Using the Ubuntu Update Manager." For information about updating a Kubuntu system, see the section of Chapter 23 entitled "Using the Adept Updater."

Test-Driving Ubuntu or Kubuntu Linux

Figure 2-4 shows the standard Ubuntu desktop displayed after you boot from the Ubuntu Desktop DVD that is provided with this book. Figure 2-5 shows the KDE 3-based Kubuntu desktop displayed after you boot from a Kubuntu 8.04 (or earlier) Desktop CD. Figure 2-6 shows the KDE 4-based Kubuntu desktop displayed after you boot from a Kubuntu 8.10 (or later) Desktop CD. From any of these desktops, you can experiment with any of the applications provided on the Desktop CDs, including audio, video, and network-oriented applications.

Although you can experiment on your own with any of the applications that are provided on the Desktop CD, the Ubuntu Desktop CD also provides a folder of sample documents, graphics files, and so on, that simplifies experimenting with different types of applications by giving you some actual documents to experiment with. The Examples entry on the desktop of an Ubuntu Desktop CD is a symbolic link to the `/usr/share/example-content` directory.

Desktop CDs and DVDs are complete Ubuntu or Kubuntu systems, and you can follow the instructions in any of the chapters of this book to experiment with Ubuntu or Kubuntu without ever installing either to your hard drive. The down side of this type of experimentation is that any files that you create or software that you configure will usually be lost when you reboot your system.

CAUTION **For some reason, the example folder is not included on Kubuntu Desktop CDs. If you're using Kubuntu and would like to have some examples to experiment with, I've put a tar archive of this folder on my Web site at** `www.vonhagen.org/ubuntu/u810-example-content.tgz`**. After downloading this file and saving it to your Kubuntu system, you can unpack its contents using the following command from a Konsole window:**

```
tar zxvf u810-example-content.tgz
```

You can then change directory to the `example-content` **directory and experiment with these files. Remember that nothing that you download to a system that you've booted from Desktop CD is saved across reboots, so you may want to copy this to another partition or system unless you want to re-download it each time you reboot.**

This section discusses the contents of the Examples folder provided on the Ubuntu Desktop CD and the applications associated with each of the files that it contains. It also explains a variety of ways that you can save any data that you create, ranging from manually saving it to an existing filesystem or networked system to using a USB stick to create a truly portable, persistent system based on the Ubuntu or Kubuntu Desktop CDs.

Exploring the Desktop CD's Examples Folder

The Ubuntu Desktop DVD that is provided with this book includes a folder of sample documents, graphics files, and so on, that makes it easy for you to experiment with different Linux applications by giving you some actual files to open, edit, save, and generally play with. (Kubuntu users, don't despair — see the note in the previous section for information about how to download an archive file containing this directory to your machine.) The example files not only serve to demonstrate various applications, but also provide you with a convenient source of marketing literature for Ubuntu and Kubuntu. You should feel free to circulate documents such as the PDF case studies to your non-Ubuntu friends.

The documents and other files in the Examples folder are as follows:

- `case_Contact.pdf`: A case study of the adoption of Ubuntu by Contact Air Gmbh, a division of the German airline, Lufthansa. Double-clicking this item opens it in the Evince PDF viewer on an Ubuntu system, in the KPDF PDF viewer on a Kubuntu 8.04 or earlier system, or in the Okular PDF viewer on a Kubuntu 8.10 or better system.
- `case_howard_county_library.pdf`: A case study of the adoption of Ubuntu by the Howard County Library. Double-clicking this item opens it in the Evince PDF viewer on an Ubuntu system, in the KPDF PDF viewer on a Kubuntu 8.04 or earlier system, or in the Okular PDF viewer on a Kubuntu 8.10 or better system.
- `case_KRUU.pdf`: A case study of the adoption of Ubuntu by KRUU, a non-profit radio station. Double-clicking this item opens it in the Evince PDF viewer on an Ubuntu system, in the KPDF PDF viewer on a Kubuntu 8.04 or earlier system, or in the Okular PDF viewer on a Kubuntu 8.10 or better system.
- `case_OaklandUniversity.pdf`: A case study of the adoption of Ubuntu by the University of Oakland. Double-clicking this item opens it in the Evince PDF viewer on an Ubuntu system, in the KPDF PDF viewer on a Kubuntu 8.04 or earlier system, or in the Okular PDF viewer on a Kubuntu 8.10 or better system.
- `case_oxford_archaeology.pdf`: A case study of the adoption of Ubuntu by Oxford Archaeology, and English excavation and heritage management firm. Double-clicking this item opens it in the Evince PDF viewer on an Ubuntu system, in the KPDF PDF viewer on a Kubuntu 8.04 or earlier system, or in the Okular PDF viewer on a Kubuntu 8.10 or better system.
- `case_Skegness.pdf`: A case study of the adoption of Ubuntu by the Skegness Grammar School. Double-clicking this item opens it in the Evince PDF viewer on an Ubuntu system, in the KPDF PDF viewer on a Kubuntu 8.04 or earlier system, or in the Okular PDF viewer on a Kubuntu 8.10 or better system.
- `case_ubuntu_johnshopkins_v2.pdf`: A case study of the adoption of Ubuntu by Johns Hopkins University. Double-clicking this item opens it in the Evince PDF viewer on an Ubuntu system, in the KPDF PDF viewer on a Kubuntu 8.04 or earlier system, or in the Okular PDF viewer on a Kubuntu 8.10 or better system.
- `case_ubuntu_locatrix_v1.pdf`: A case study of the adoption of Ubuntu by Locatrix Communications. Double-clicking this item opens it in the Evince PDF viewer on an Ubuntu system, in the KPDF PDF viewer on a Kubuntu 8.04 or earlier system, or in the Okular PDF viewer on a Kubuntu 8.10 or better system.
- `case_Wellcome.pdf`: A case study of the adoption of Ubuntu by the Wellcome Trust Sanger Institute. Double-clicking this item opens it in the Evince PDF viewer on an Ubuntu system, in the KPDF PDF viewer on a Kubuntu 8.04 or earlier system, or in the Okular PDF viewer on a Kubuntu 8.10 or better system.

- `fables-01_01_aesop.spx`: A sample Speex audio file (an open source audio compression format) in Ogg container format. Double-clicking this item opens and plays it using the Totem video player application on an Ubuntu system, or using the Amarok audio player on Kubuntu systems.
- `gimp-ubuntu-splash.xcf`: A sample Ubuntu splash screen in the native open source XCF format used by the GIMP imagine manipulation program. Double-clicking this item opens it in the GIMP on an Ubuntu system, or in Okular on a Kubuntu 8.10 or better system. There is no default file associate for XCF files on Kubuntu 8.04 and earlier systems, though the file can be manually opened in Gwenview.
- `kubuntu-leaflet.png`: A sample Kubuntu document in the Portable Network Graphics (PNG) format. Double-clicking this item opens it in the GIMP on an Ubuntu system, in Gwenview on a Kubuntu 8.04 and earlier system when using Dolphin, in Konqueror on a Kubuntu 8.04 and earlier system when using Konqueror, and in Okular on a Kubuntu 8.10 and later system.
- `logos`: A folder containing sample Ubuntu and Kubuntu logos in PNG format. Double-clicking either of these items opens it in the GIMP on an Ubuntu system, in Gwenview on a Kubuntu 8.04 and earlier system when using Dolphin, in Konqueror on a Kubuntu 8.04 and earlier system when using Konqueror, and in Okular on a Kubuntu 8.10 and later system.
- `oo-about-these-files.odt`: A sample OpenOffice.org document that provides information about the files in the Examples folder. Double-clicking this item opens it in OpenOffice Writer.
- `oo-about-ubuntu-ru.rtf`: A sample Russian document in Microsoft's Rich Text Format (RTF) markup. Double-clicking this item opens it in OpenOffice Writer.
- `oo-derivatives.doc`: A sample document in Microsoft Office/Word document format that discusses derivatives of the Ubuntu Linux distribution, such as Edubuntu, Kubuntu, and Xubuntu. Double-clicking this item opens it in OpenOffice Writer.
- `oo-maxwell.odt`: A sample OpenOffice.org document that shows the mathematical capabilities of OpenOffice.org Writer by discussing Maxwell's equations. Double-clicking this item opens it in OpenOffice Writer.
- `oo-payment-schedule.ods`: A sample OpenOffice.org spreadsheet that shows calculations and other capabilities. Double-clicking this item opens it in OpenOffice Calc.
- `oo-presenting-kubuntu.odp`: A sample OpenOffice.org presentation that discusses Kubuntu. Double-clicking this item opens it in OpenOffice Impress.
- `oo-presenting-ubuntu.odp`: A sample OpenOffice.org presentation that discusses Ubuntu. Double-clicking this item opens it in OpenOffice Impress.
- `oo-trig.xls`: A sample Microsoft Office/Excel document that shows calculations and the graphical display of results. Double-clicking this item opens it in OpenOffice Calc.
- `oo-welcome.odt`: A sample OpenOffice.org document that contains embedded graphics and provides an introduction to Ubuntu Linux. Double-clicking this item opens it in OpenOffice Writer.
- `Ubuntu_Free_Culture_Showcase`: A folder containing the winning audio and video files from the Ubuntu 8.10 Free Culture Showcase (`https://wiki.ubuntu.com/UbuntuFreeCultureShowcase`), which sought media files that were released under a Creative Commons Attribution ShareAlike license (`https://wiki.ubuntu.com/ShareAlike`) to use when demonstrating (i.e., showing off) a new Ubuntu system. Double-clicking the OGA file in this directory plays it in Totem on an Ubuntu system, and in Amarok on Kubuntu systems. Double-clicking the OGV file in this directory opens it in Totem on an Ubuntu system, in Dragon player on a Kubuntu 8.10 or newer system, or in Kaffeine on a Kubuntu 8.04 or earlier system.

Double-clicking on these sample documents and files (or single-clicking them on a default Kubuntu system) gives you the opportunity to experiment with the associated software package without having to start from scratch. You cannot modify these documents in place because they are loaded from the CD. To modify any of them, copy the document or file to the login directory and make it writable. For information about changing Linux file permissions, see the section of Chapter 4 entitled "Understanding Linux Permissions."

Accessing Your Hard Drive from the Desktop CD

Booting a desktop version of Ubuntu or Kubuntu on your Microsoft Windows or other Linux system is fun, but if you do any real work while in "Desktop CD mode," you may actually want to save it for future reuse. Because the Ubuntu and Kubuntu Desktop CDs run from an in-memory filesystem (known as a RAM disk), the key to saving data that you've created when running Ubuntu or Kubuntu from a Desktop CD is being able to access the existing filesystems on your hard drive(s) and saving your data there, or using an external persistent storage device as explained later in this chapter in the section entitled "Using Desktop CD Persistence." To access the filesystems on your local disk drives, you have to explicitly make them available (mount them) when running Ubuntu or Kubuntu, and this requires that you use a few special utilities — and, as always, you have your choice of graphical or command-line utilities.

The next sections explain how to mount and access your existing filesystems from different types of Desktop CDs so that you can preserve files that you've created when running Ubuntu or Kubuntu from a Desktop CD.

Accessing Existing Partitions from a Desktop CD

Microsoft Windows systems use either the NTFS (NT File System) or the older FAT32 (32-bit File Allocation Table) filesystems. Luckily, Linux provides complete support for both of these types of filesystems, so once you've booted from an Ubuntu or Kubuntu Desktop CD, it's easy to mount the filesystems that are already present on your system's hard drive so that you can access your existing files or save any files that you have created while running from a Desktop CD.

Accessing Existing Partitions from an Ubuntu Desktop CD Using Nautilus

If you've booted from an Ubuntu Desktop CD, you can do the following to access the existing filesystems on your system's hard disk(s):

1. After booting from the Desktop CD, select the Places -> Computer menu item. A Nautilus File Manager window displays, showing all of the available disk partitions in the right-hand pane and also listing them in the Places list in the left-hand pane, as shown in Figure 2-19.

 Figure 2-19 shows mountable partitions that are 6.3GB, 21.0GB, and 60GB in size. The CD/DVD drive contains the Ubuntu Desktop CD, and the Filesystem entry represents the in-memory RAM disk that contains the root filesystem that is created when you boot from a Desktop CD.

2. To mount any of the partitions from your hard disk, left-click on the name of the partition in the Places list in the left pane. An icon for the partition displays on your desktop, and the contents of that partition display in the right pane of the Nautilus window, as shown in Figure 2-20.

You should be able to identify the partition and type of partition that you've mounted by examining the files that it contains. For example, the partition shown in Figure 2-20 is clearly a Microsoft Windows partition given that it contains directories such as `WINDOWS`, `Program Files`, and `My Music`.

FIGURE 2-19

Available disk partitions in Nautilus

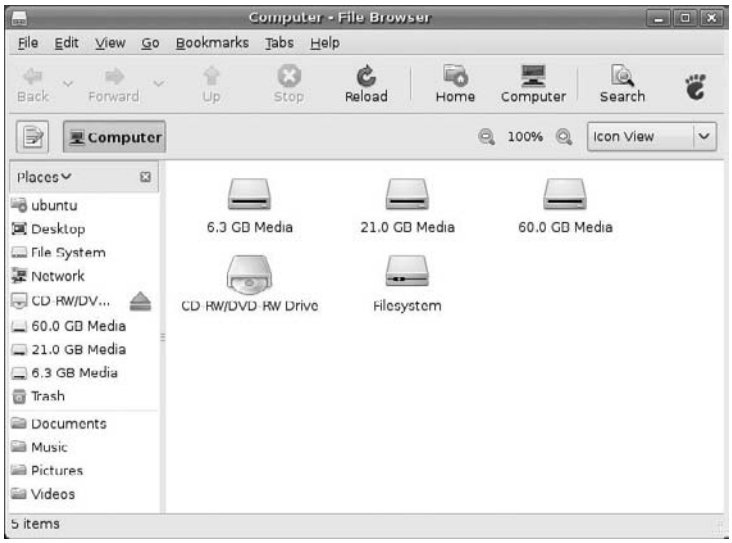
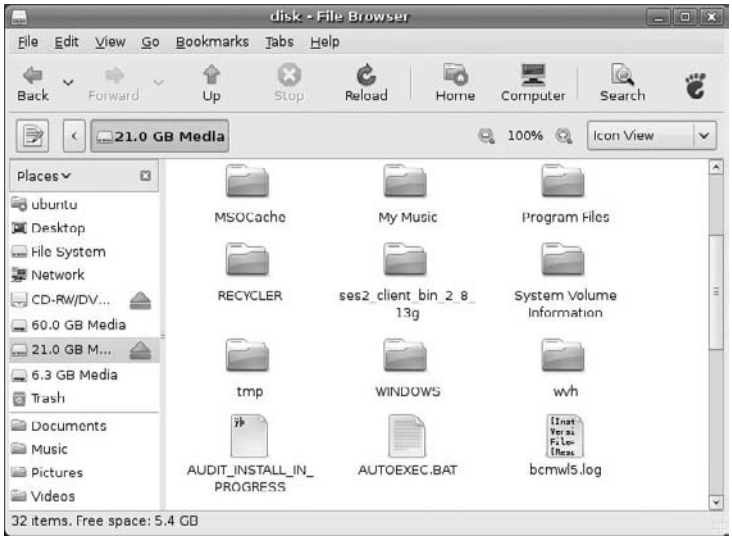


FIGURE 2-20

Viewing a Windows partition in Nautilus



Once you've mounted a partition, a graphical arrow displays to the right of its name in the Places list, showing that it is currently mounted. After dragging and dropping any files or directories that you want to save from your Live CD session to that partition in Nautilus, you can left-click that arrow to unmount the partition. As a general rule, it's a good idea to unmount any of your hard disk partitions that you've mounted before shutting down or rebooting from the Live CD.

Accessing Existing Partitions from a Kubuntu Desktop CD Using Dolphin

If you've booted from a Kubuntu 8.04 or later Desktop CD, you can do the following to access the existing filesystems on your system's hard disk(s):

1. After booting from the Desktop CD, right-click on the desktop and select the Run Command menu item. Enter the command `dolphin` and press Enter/Return on your keyboard. A Dolphin File Manager window displays.
2. If you are using a Kubuntu 8.04 Desktop CD, select the Storage Media entry in the left-hand pane of the Dolphin window. Icons for each available partition display in the center pane, as shown in Figure 2-21, which shows mountable partitions that are 6.3GB, 20GB, and 60GB in size.
3. If you are using a Kubuntu 8.10 or later Desktop CD, the available partitions will be listed as volumes in the left-hand pane of the Dolphin window, as shown in Figure 2-22, which lists three available volumes. Although the version of Dolphin provided with Kubuntu 8.10 doesn't identify the size of the volumes, it does identify the type of filesystem that each contains, which makes it easy to locate the Windows partitions.

FIGURE 2-21

Available disk partitions in Dolphin on a Kubuntu 8.04 System

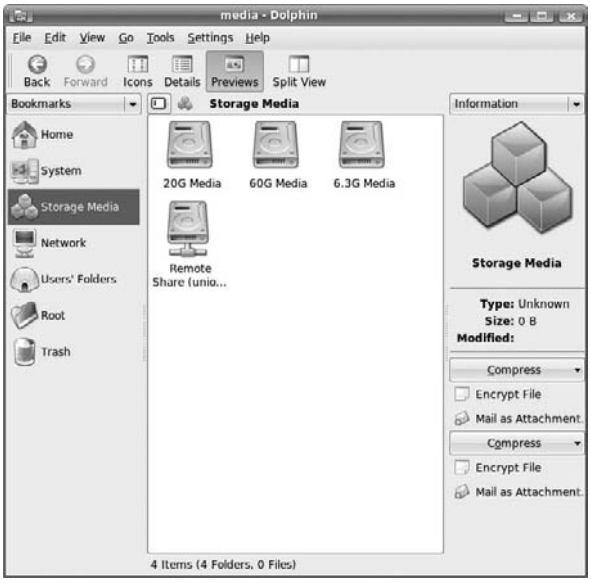
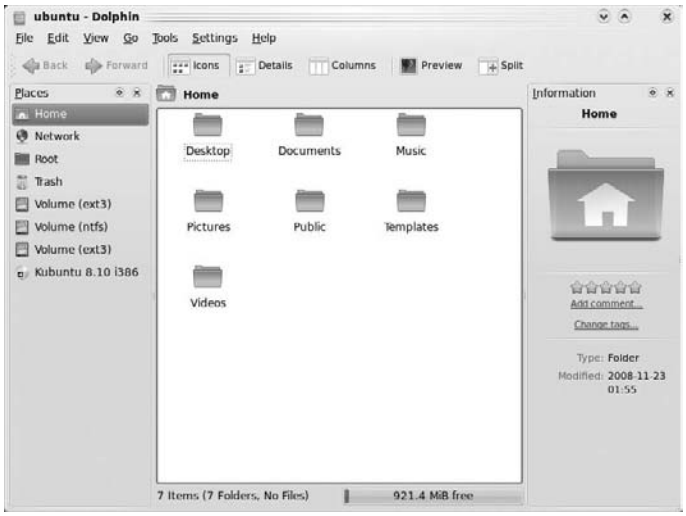


FIGURE 2-22

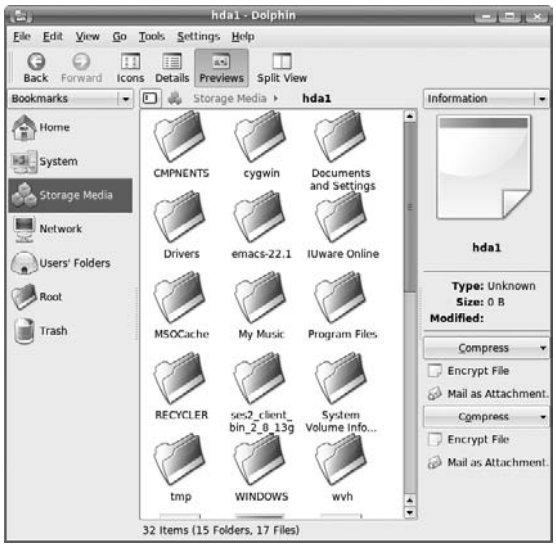
Available volumes in Dolphin on a Kubuntu 8.04 System



4. If you are using a Kubuntu 8.04 system, left-click on any of the media icons in the central Dolphin pane to mount that partition and display its contents, which will look something like Figure 2-23.

FIGURE 2-23

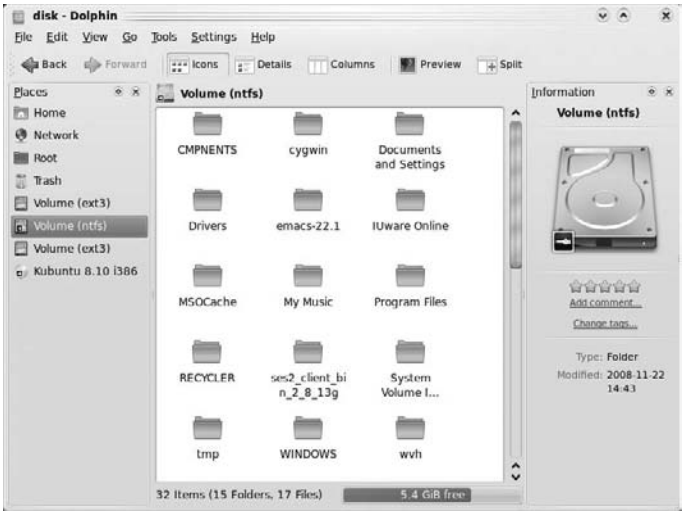
A Windows partition in Dolphin on Kubuntu 8.04



5. If you are using a Kubuntu 8.10 or better system, left-click on any of the volume icons in the left-hand Dolphin pane to mount that volume and display its contents in the central Dolphin pane, which will look something like Figure 2-24.

FIGURE 2-24

A Windows partition in Dolphin on Kubuntu 8.10 or better



You can now drag and drop files from your desktop to and from your mounted Windows or Linux partitions. Once you're done, you can restart your system, and you will find the files and directories that you created in the partition just as if you had created them there from your system's native Windows or Linux operating system.

Accessing Existing Partitions from the Command Line

The previous two sections explained how to use the standard file managers on Ubuntu and Kubuntu systems to locate and mount partitions from your system's hard drive when you've booted from an Ubuntu or Kubuntu Desktop CD. Many people prefer to work from the command line rather than from a file manager whenever possible (hmm, my hand seems to have raised itself), so this section summarizes how to locate the partitions on your system's hard drive, identify the type of filesystem that each contains, and mount those partitions so that you can copy data to and from them, all from the comfort of the Linux command line.

To locate, mount, use, and unmount partitions from the command line, do the following:

1. If you are using an Ubuntu Desktop CD, select the Applications -> Accessories -> Terminal menu command to start a GNOME terminal. If you are using a Kubuntu Desktop CD, right-click on the desktop, select Run Command from the menu that displays, enter **konsole**, and press Enter/Return to start a KDE Konsole window.
2. Enter the `sudo -s` command to become the root user.

3. Use the `egrep` command, which searches for regular expression patterns in a specified file, to identify the disks and partitions that were found when you booted your system, as in the following example:

```
# egrep '[hs]d[a-z][0-9]' /var/log/messages
Nov 22 19:36:43 ubuntu kernel: sda1 sda2 sda3 sda4 < sda5 sda6 >
Nov 22 19:36:43 ubuntu kernel: Adding 4000176k swap on /dev/sda3
...
```

The first of the output lines shows that a hard drive named `/dev/sda` was located during the boot process, and that it contains six partitions, the last two of which are logical partitions within an extended partition.

4. Use the `fdisk -l` command to get more detailed information about the partitions on the disk, as shown in the following example. The last column identifies the type of partition associated with each Linux device name:

```
# fdisk -l /dev/sda

Disk /dev/sda: 160.0 GB, 160041885696 bytes
255 heads, 63 sectors/track, 19457 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Disk identifier: 0xedaaedaa

   Device Boot  Start      End     Blocks   Id  System
/dev/sda1   *     766     3315   20482875    7  HPFS/NTFS
/dev/sda2           1      765    6144831   83  Linux
/dev/sda3        3316     3813    4000185   82  Linux swap
/dev/sda4        3814    19457  125660430    5  Extended
/dev/sda5        3814    11108   58597056   83  Linux
/dev/sda6       11109    19457   67063311   83  Linux
```

5. Mount the partition that contains the files that you want to access, or to which you want to save files that you've created using the Desktop CD. For example, to mount the NTFS partition `/dev/sda1` on the Linux directory `/mnt`, you would use a command like the following:

```
mount -t ntfs-3g /dev/sda1 /mnt
```

If `/dev/sda1` had been an old-style FAT32 partition, you would use a command like the following to mount it:

```
mount -t vfat /dev/sda1 /mnt
```

Similarly, if you wanted to mount `/dev/sda5`, which is one of your Linux partitions, you would use a command like the following (assuming that the filesystem used the standard ext3 filesystem format):

```
mount -t ext3 /dev/sda5 /mnt
```

As explained in more detail in Chapter 4, Linux systems access partitions by mounting them on existing directories in the filesystem. The `/mnt` directory is a standard directory that is provided on all Linux systems for temporary mount purposes. Note that you can mount only one filesystem on a given directory — you could execute only one of the previous commands.

6. You can now copy files to and from the directory `/mnt`. When you're done copying files to this partition or working with the files that it contains, you should unmount it using the following command:

```
umount /mnt
```

7. Drop the permissions of the root user by typing `exit` or Ctrl+D.

Congratulations — you are now an official Linux geek! See Chapter 4 for more detailed information about Linux concepts such as mounting and unmounting filesystems.

Copying Files to Other Machines Over a Network

If you create or modify files while running from a Desktop CD, want to save that work, but don't want to save it to any of your system's hard disk partitions (as described earlier in this chapter), you can still save those files by copying them over the network to another machine.

Linux provides several ways to enable copying files to other systems over the network. Some of the most common and popular of these are as follows:

- The `ftp` (file transfer protocol) utility, which establishes a connection to another machine and enables you to interactively transfer files in either direction
- The `scp` (secure copy) utility, which enables you to quickly transfer files to other machines over a secure, encrypted connection

These commands are explained in detail in Bonus Chapter 1, "Connecting to Other Systems." They're summarized here so that you know that there are ways to save any work that you've done from an Ubuntu or Kubuntu Desktop CD that don't involve local storage. This can be very useful if, for example, you are using an Ubuntu or Kubuntu Desktop CD to recover data from a crashed system or one that can't boot on its own for some reason.

Using Desktop CD Persistence

Live CDs such as the Ubuntu and Kubuntu Desktop CDs provide an easy way to temporarily turn any compatible computer system into an Ubuntu or Kubuntu Linux system. However, the fact that most Live CDs don't provide any sort of persistent storage for configuration data, custom applications, and so on (other than manually copying files to other local or networked storage) traditionally limited the practical usability of Live CDs to simple demonstrations, one-time experimentation, and general system repair and recovery efforts. Luckily, things are better now — thanks to some really clever people, it's fairly easy to automatically save changes that you make when running from the equivalent of an Ubuntu or Kubuntu Desktop CD — and without having to deal with manually mounting and unmounting filesystems, saving your work there, and so on.

Earlier Ubuntu releases introduced a mechanism for associating persistent storage with the Live CD boot process. This mechanism enabled you to format an external storage device with a special name (caspar-rw)

and use that at boot time (in conjunction with the Desktop CD), adding a special kernel boot option (persistent) to tell the boot process to use the contents of that device to restore desktop settings, personal data, and even installed applications. This approach had a few problems, the most significant of which was that it required both an external storage device (typically a USB flash memory stick) and an Ubuntu or Kubuntu Desktop CD. Aside from the hassle of carrying around both of these, my particular complaint was that my Desktop CDs kept breaking when I put them in my back pocket and then sat on them.

Luckily, improvements in flash drive size, price reductions for flash drives, and software improvements have made being able to do real work from an Ubuntu or Kubuntu Desktop CD even easier. Ubuntu 8.10 introduces an easy-to-use command for creating a bootable USB flash drive (generally referred to as a *stick*) that works like the Desktop CD but provides persistent storage when booting using the Try Ubuntu without any change to your computer option. This gives you a portable Ubuntu system from which you can either run Ubuntu or install it on another computer system without requiring the use of a CD drive. All you need nowadays for this type of portable Ubuntu system is a USB flash drive that is at least 2GB in size. Unfortunately, an equivalent application for creating bootable USB flash drives isn't provided on the Kubuntu 8.04 or later Desktop CDs (yet). However, don't worry Kubuntu fans — I'll also explain how to use an Ubuntu Desktop CD to create a bootable USB flash drive using an ISO image of your favorite Kubuntu Desktop CD.

CAUTION **If you want to create a bootable USB flash drive manually, the folks at PenDriveLinux.com (`www.pendrivelinux.com`) have some truly useful tutorials that will walk you through the process of manually creating a bootable Ubuntu 8.10 system on a USB stick, a bootable Xubuntu system on a USB stick, a bootable Kubuntu 8.10 system on a USB stick from Windows, and just about anything else that you could conceivably boot from a USB stick. Their tutorials are clearly explained and eminently readable, and also often include scripts and configuration files that you can download from their site to make things even easier for you.**

Flash drives are great — they are fast, use far less energy than hard drives, don't generate heat, don't make noise, and have no moving parts that can stop moving when you drop them down a flight of stairs. On the down side, they do have one potential limitation that you'll hear people talk about — you can only write to each underlying storage unit a certain number of times. (Flash storage units are usually referred to as *cells* or *address blocks*, and are analogous to the sectors on a standard hard drive.) Luckily, most flash drives use a technique known as *wear-leveling* to evenly distribute writes across the drive, so that no single area on the drive gets written to more frequently than another. The longevity of a modern flash drive is usually predicted to be at least 5 or 10 years. If you're still running Ubuntu or Kubuntu 8.10 off the same flash drive five years from now, it might be time to finally consider an upgrade. You are backing up your "real work" anyway, right?

NOTE **In order to use the bootable USB flash drives discussed in the next two sections, your computer system must be able to boot from a USB device. Most modern PCs support this, but you should certainly check before bothering to create a bootable USB stick that you can't use. Booting from USB devices is usually a standard BIOS boot option, but some systems also require that you manually enable USB BIOS support before you can boot from a USB device.**

Creating a Bootable Ubuntu USB Stick

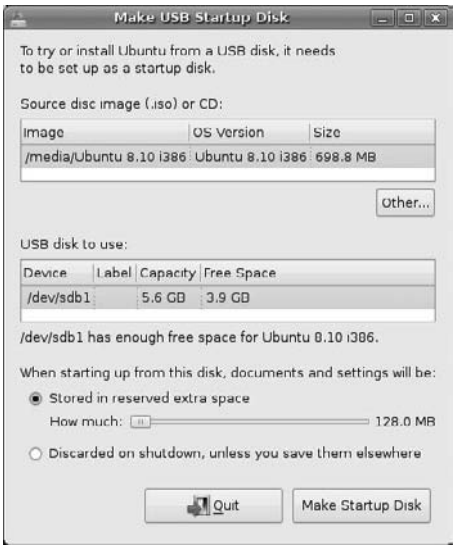
Ubuntu 8.10 provides an easy, graphical mechanism for creating a bootable USB flash drive that works much like the Desktop CD, but which also provides persistent storage so that any system settings that you configure, applications that you install, and files that you create are preserved across reboots. This application will try to reuse any FAT32 partition that you have on your USB stick, so if you have formatted your USB stick to use another type of filesystem or have created multiple partitions on your USB stick, you may want to see the sidebar on reformatting a USB flash drive, later in this chapter.

To use this application, do the following after booting from an Ubuntu Desktop CD using the "Try Ubuntu without any change to your computer" option:

1. Select the System -> Administration -> Create a USB startup disk menu option. A dialog like the one shown in Figure 2-25 displays.

FIGURE 2-25

Creating a USB startup disk on Ubuntu 8.10



2. Verify that the mount point for your Ubuntu Desktop CD is listed and highlighted in the Source disk image (.iso) or CD area at the top of the dialog. If not, click Other and browse to the correct location.
3. Verify that the USB disk partition on which you want to create a bootable Ubuntu system is listed and highlighted in the USB disk to use area in the middle of the dialog. The line below the USB disk list identifies whether the selected USB partition has sufficient space for installing a bootable system.
4. To create a bootable USB stick that supports the persistent storage of applications that you install, files and directories that you create, and applications and system configuration information, make sure that the Stored in reserved extra space radio button is selected at the bottom of the dialog shown in Figure 2-25. Drag the slider below this radio button to identify the amount of space that you want to allocate on the selected partition for storing persistent data.

 If you do not want to create any persistent storage, click the Discarded on shutdown . . . radio button. If you select this radio button, your bootable USB stick will behave exactly like an Ubuntu Desktop CD. Any files or directories that you create, applications that you install, and application and system configuration information that you modify will not be preserved when you shut down a system that you have booted from the USB stick.

Reformatting a USB Flash Drive

The Ubuntu Create a USB startup disk application will try to reuse any existing FAT32 (VFAT) partition on your USB stick. If you've reformatted the USB stick to contain multiple partitions, changed the filesystem type of any of your existing partitions, or already have multiple partitions on the stick but don't need to use them, this sidebar explains how to reformat your flash drive so that it contains a single FAT32 partition that the Create a USB startup disk option can make the most of. As you might expect, this will delete any information that is currently present on the USB stick, so make sure that you've copied anything that you ever want to see again elsewhere before following the instructions in this sidebar.

To reformat a USB stick, do the following:

1. Select the Applications -> Accessories -> Terminal menu command to start a GNOME terminal.
2. Enter the `sudo -s` command to become the root user.
3. Use the `df` command to see the mounted partitions on your system. If any of the partitions from your USB stick are mounted, unmount each of them using the `umount` command.
4. Execute the `fdisk` command with the name of the device corresponding to your USB stick as an argument. You must specify the name of the device, not a partition on that device. For example, if your USB stick were `/dev/sdb`, you would execute the command `fdisk /dev/sdb`.
5. Type `p` and press Enter/Return to see a list of all of the partitions on the USB stick. Use the `d` *`number`* command to delete each of these, where *number* is the number of the partition as listed in the device names in the first column of the output displayed by the `p` command. When you're done, type the `p` command again to list the partition table — you should not see any partitions listed.
6. Use the `n` command to begin creating a new partition on the flash drive. Enter `p` to create this as a primary partition, and enter `1` as the partition number. Press Return when prompted for the starting block of the new partition to accept the default value of 1. When prompted for the end of the partition, press Enter/Return again to use the entire disk.
7. Type `t` and press Enter/Return to begin changing the type of the first partition on the drive. Enter `1` to identify the partition that you want to change, and enter `c` (Windows 95 FAT32 LBA) as the type of the partition.
8. Type `a` and press Enter/Return to set your new partition as *active* (i.e., bootable). When prompted for the partition number, enter `1` and press Enter/Return.
9. Type `w` to write your updated partition table to the USB stick. Type `q` to exit from the fdisk application.

After repartitioning your USB stick, you must remove and reattach it to your system to cause the system to read the new partition table. After re-attaching the drive, use the `dmesg` command to display the last few kernel messages, one of which should identify the device name and available partitions on the USB stick that you just reattached. You can use the `fdisk -l /dev/`*`DEVICE`* command to list the partitions on the device named DEVICE to verify that you know the device name of the re-attached drive — it is not necessarily the same device name that it had before. Once you are absolutely sure that you know the device name of the re-attached USB stick, you can format it using a command like the following:

```
mkfs.vfat -F 32 -n NAME /dev/PARTITION
```

Replace *`NAME`* with the name that you want to assign to your USB stick, and replace *`PARTITION`* with the name of the partition on the device, which would be something like sdb1, sdc1, and so on. You're now ready to create a bootable Ubuntu or Kubuntu system on your USB stick.

5. Click Make Startup Disk to begin creating the bootable USB stick. This process takes a few minutes. A dialog displays when the bootable system has been successfully created in the selected USB device partition. Click OK to close this dialog and the startup disk creation application.

At this point, you can remove the USB stick from your system and test it on another system, or you can reboot your current system to test rebooting from the USB stick. Make sure that you have activated booting from USB devices in the BIOS of any system that you use to test the USB stick, and that the USB boot device has a higher priority than your system's CD, DVD, or hard drives.

Creating a Bootable Kubuntu USB Stick

As mentioned at the beginning of this section, only the Ubuntu Desktop CDs provide the Create a USB startup disk application. This is somewhat sad, because "Kubuntu on a stick" is just as attractive as "Ubuntu on a stick." Thanks to the flexibility of the USB Creator application, however, you can easily create a bootable Kubuntu 8.10 stick.

To create a bootable Kubuntu stick, boot from an Ubuntu Desktop CD and download or transfer a Kubuntu 8.10 ISO image to the live system one it is up and running. Follow the instructions in the "Creating a Bootable Ubuntu USB Stick" section, but click Other… below the Source disk image (.iso) or CD section, browse to the location of the Kubuntu 8.10 ISO image, select it, and click OK. When the USB Creator application redisplays, select that ISO as the Source disk. You can then follow the rest of the steps described in the previous section. When the application finishes, you will have a bootable Kubuntu 8.10 USB stick. Not too shabby!

Wubi: Installing Ubuntu or Kubuntu on Windows

Versions of the Ubuntu and Kubuntu Desktop CDs prior to version 8.04 provided a variety of open source software packages that were compiled for Microsoft Windows systems, enabling Windows users to try out some of the amazing software that they could see and use after booting from the Desktop CDs in Live CD mode. These packages included Mozilla Firefox, Mozilla Thunderbird, the GIMP image creation and editing package, and several others.

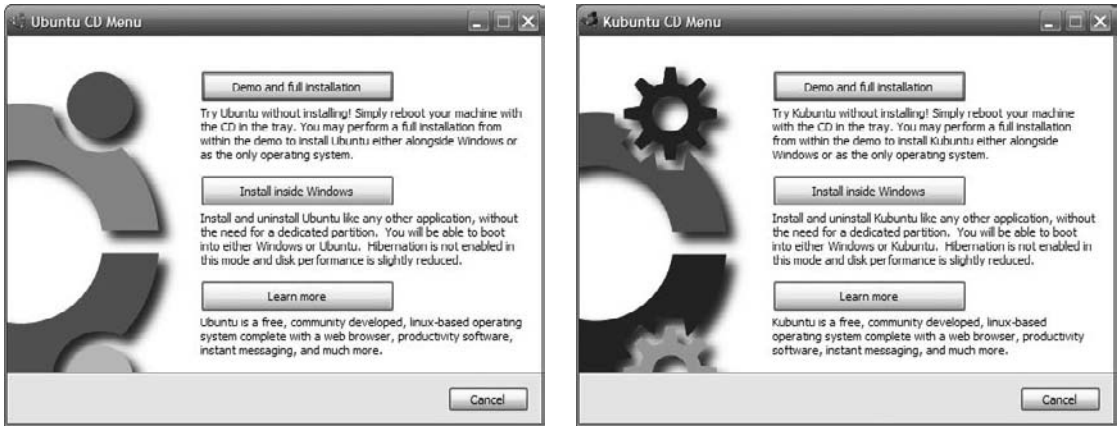
Versions 8.04 and later of the Ubuntu and Kubuntu Desktop CDs stopped offering the Windows versions of these packages, providing an even better option for Windows users: the option to install Ubuntu or Kubuntu onto any existing Windows system without repartitioning your existing disk, adding a boot option to your Windows system that enables you to boot into either Ubuntu/Kubuntu or your existing Windows system. This capability is provided by an impressive package called Wubi (Windows-based Ubuntu Installer — `http://wubi-installer.org`). Installing Ubuntu or Kubuntu onto an existing Windows system requires a minimum of 4GB of disk space (5GB or better is recommended), but gives you a complete Ubuntu or Kubuntu system that is installed into a new directory on your Windows system. On the Windows side, a Wubi-based installation looks like any other Windows application, and can even be removed using the traditional Add/Remove Applications Control Panel application.

A Wubi installation works by creating an Ubuntu directory at the top level of your system's hard drive and populating that with the files that Wubi requires to boot Ubuntu or Kubuntu, including a filesystem image for that distribution. (This directory is named Ubuntu regardless of whether you are installing Ubuntu or Kubuntu.) The installation process adds an Ubuntu or Kubuntu option to your Windows boot menu, and then installs the boot loader (`wubildr` — the Wubi loader) that it uses to mount and access the filesystem image as the application that is executed when the Ubuntu or Kubuntu boot option is selected.

To take advantage of this installation option, insert an Ubuntu or Kubuntu 8.04 or better Desktop CD into the CD/DVD drive of an existing Windows system. Depending on whether you are using an Ubuntu or Kubuntu Desktop CD, one of the dialogs shown in Figure 2-26 displays.

FIGURE 2-26

Ubuntu and Kubuntu CD dialogs on a Microsoft Windows system



To install Ubuntu or Kubuntu on your existing Windows system, click Install Inside Windows. Depending on whether you are using an Ubuntu or Kubuntu Desktop CD, one of the dialogs shown in Figure 2-27 displays.

By default, the Username field shows the name of the Windows user that is currently logged in. If you want to use a different user when logging into your Wubi-based Ubuntu or Kubuntu installation, enter the username that you want to use.

Next, enter the password that you want this user to have on your Wubi-based Ubuntu or Kubuntu installation, verifying that you haven't made a typo by entering the password twice.

Finally, if you have any special Accessibility requirements, click Accessibility to display a dialog like the one shown in Figure 2-28, which shows the Ubuntu version of this dialog. After setting any accessibility options that you want to use, click Next to return to the main installation dialog.

To begin the installation process, click Install. You will see a progress dialog as Wubi verifies the checksums of the files and directories that it is installing, installs Ubuntu or Kubuntu into the filesystem image, and so on. A summary dialog displays when the installation completes, offering you the option to reboot now or later to start using your new Wubi-based Ubuntu or Kubuntu installation.

After installing Ubuntu or Kubuntu and Wubi, you will see a screen like that shown in Figure 2-29 the next time you boot your Windows system.

This screen shows the boot screen for a Wubi-based Ubuntu installation — a Wubi-based Kubuntu installation would offer a Kubuntu option. To boot into your existing Windows installation, simply press Enter/Return on your keyboard. To boot into Ubuntu (or Kubuntu), use the arrow keys to select that boot option and press Enter/Return.

FIGURE 2-27

Beginning the Ubuntu or Kubuntu installation process on Windows



FIGURE 2-28

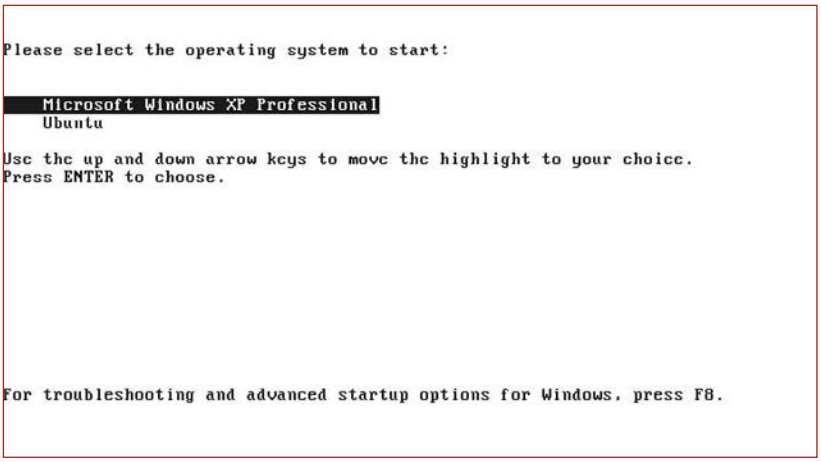
Accessibility settings for a Wubi-based installation on Windows



Wubi is another shining example of the creativity of the open source community; it provides yet another way for Windows users to experience the power of Ubuntu and Kubuntu Linux.

FIGURE 2-29

A Wubi boot screen on Windows



Summary

This chapter explained how you can take Ubuntu or Kubuntu Linux for a test drive or even permanently install them after booting from the Ubuntu or Kubuntu Desktop CD or DVD that is appropriate for the type of personal computer that you're currently using. (This book includes an Ubuntu Desktop DVD for x86 systems.) This chapter also explained where to get the Kubuntu Desktop CDs if you're a KDE fan — and which version of Kubuntu to select, depending on your requirements and curiosity. Although I hope that you've installed Ubuntu or Kubuntu Linux permanently on your computer system, this chapter also discussed various ways of experimenting with them and saving your data if you're not yet ready to install one of them permanently on your computer.

If you're not ready to completely commit yourself to Ubuntu or Kubuntu Linux, the Ubuntu and Kubuntu Desktop CDs also enable you to install them on an existing Windows system as a dual-boot system, or within an existing Windows partition without re-partitioning or making any significant changes to your existing Windows system (aside from adding a boot option and using up some disk space, that is). The latest generation of the Ubuntu and Kubuntu Desktop CDs make it easier than ever before to experience the power, flexibility, and freedom provided by the Ubuntu and Kubuntu Linux distributions.

The next chapter discusses other ways of installing Ubuntu or Kubuntu Linux using CDs other than the Desktop CDs that this chapter focused on. These include CDs that enable you to install Ubuntu or Kubuntu on systems with less than 512MB of memory or in OEM scenarios by using a text-based installer, and the Ubuntu server CD, which makes it easy to install Ubuntu for server systems and provides some default configurations to help you get up and running with Ubuntu as a core component of your SOHO or enterprise infrastructure.

Chapter 3

Installing Ubuntu and Kubuntu on Special-Purpose Systems

IN THIS CHAPTER

- Getting other Ubuntu/Kubuntu CDs
- Installing Ubuntu Servers
- Resolving installation problems
- Installing from an alternate CD

The previous chapter discussed using the graphical installer that is provided as part of the Ubuntu Desktop disc, which is the DVD that is included with this book, and its Kubuntu equivalents, the Kubuntu 8.04 and Kubuntu 8.10 Desktop CDs. Using any of these discs and their installer is the standard way to install an Ubuntu or Kubuntu Desktop system — which one you decide to install is up to you. However, as discussed in the section "Ubuntu CDs" in Chapter 1, other CDs are also available from the Ubuntu and Kubuntu folks. These are the Ubuntu Server Install CD and the Ubuntu and Kubuntu Alternate Install CDs:

- The Ubuntu Server CD installs Ubuntu without a GUI for server systems and therefore avoids the whole GNOME/KDE/Ubuntu/Kubuntu issue.
- The Ubuntu and Kubuntu Alternate CDs install Ubuntu or Kubuntu using a quasi-graphical installer that enables you to install Ubuntu or Kubuntu on systems without sufficient memory to run the traditional graphical installer found on the Ubuntu and Kubuntu Desktop CDs.

This chapter discusses how to obtain these CDs and how to use the installation alternatives that each of them provides.

The Ubuntu and Kubuntu Desktop CD's graphical installers differ from the quasi-graphical, Debian-based installers used on the Server and Alternate Install CDs. Before providing a graphical installer on the Desktop CDs, the Debian-based installers used by Ubuntu and Kubuntu received their share of the criticism traditionally aimed at the Debian installer — "It's not fancy enough," "It looks like something from 1985," "It doesn't use all of the whizzy features of my five-dimensional, accelerated graphics card," and so on. These are all true, and for good reasons. Although it's nice to have a fancy, mouse-oriented, graphical installation program, an installer is something that you use once per system and then forget about. An installer therefore has to be rock-solid, easy-to-understand, and easy-to-use. It also has to work on any computer system from the most wretched, low-resolution VGA system to the high-end

gaming systems of today, which provide stunning resolution and visual nuances that many people can't even detect. Server systems often run headless (i.e., without a graphical console), and therefore an installer that works on those types of systems is fairly important.

Getting a Different Install CD

Ubuntu discs aren't available in stores, which is fine because a DVD suitable for test drives and installation is provided with this book, and the Kubuntu and other Ubuntu installation CDs are readily available over the Internet. If you have a reasonably fast Internet connection and a CD or DVD burner, the easiest way to get a different install CD or DVD than the one provided in this book is to download an ISO image for that CD. (An ISO image is a file containing an image of a CD-ROM or DVD that is in the format mandated by the International Standard Organization 9660 specification.) See the section "Ubuntu CDs" in Chapter 1 for more information about the Ubuntu and Kubuntu discs that are available other than the one provided in this book. Once you've decided which CD meets your needs, go to the page at `www.ubuntu.com/download` to get other Ubuntu CDs or `www.kubuntu.org/getkubuntu/download` to get Kubuntu CDs, select a mirror site in your country or a site that's closest to you in general, and select the Install CD that is associated with your platform and the type of system that you want to install. Save the downloaded ISO image to a file, burn that ISO image to a CD once the download completes, and you're ready to go!

If you have a slow Internet connection or do not have a CD burner, don't panic! As mentioned in the front matter for this book, the Ubuntu and Kubuntu folks will even ship you CDs, although they take a few weeks to arrive depending on your geographic location. This isn't the right solution for those of us who need instant gratification, but it may work for you. To request that the Ubuntu or Kubuntu folks ship you CDs, go to `https://shipit.ubuntu.com` or `https://shipit.kubuntu.org`, create an account by entering your e-mail address and a password, and request the Ubuntu or Kubuntu CDs for the current release. You will then have to haunt your mailbox until they actually arrive.

Booting from a Server or Alternate Install CD

Booting from an Ubuntu Server, Ubuntu Alternate Install, or Kubuntu Alternate Install CD is as easy as inserting the appropriate CD into your system's CD drive, restarting your system, and telling the system to boot from the CD drive instead of booting from a hard disk partition. On PC systems, this is done using BIOS settings.

TIP

If your x86 system boots from the hard drive even if the correct Ubuntu or Kubuntu install CD is present in your CD drive, you'll need to modify the *boot order* in your BIOS settings. The boot order is the sequence in which available devices are searched for bootable disks when your system is powered on. To do this, turn your computer on and press the key on your keyboard that gives you access to the BIOS. This is typically the Del (Delete) or F2 key for most modern systems, but the key that you'll need to press is usually identified at the bottom left of your screen when you turn on your computer and it first starts up.

Depending on the type of BIOS your computer uses, boot order settings are usually stored in an Advanced Settings or Boot screen, which you can navigate to using the arrow keys on your keyboard. Press Return to display this screen once its name is highlighted. Once this screen displays, use the down-arrow key to navigate to the First Boot Device or CD Drive entry, and see the help messages at the right side of the screen for information about how to make your CD drive the first boot device. You can then press the Escape key to exit this screen, and press F10 to save the new settings, exit the BIOS settings screen, and reboot.

Once your system begins booting from an Ubuntu Server or Alternate Install CD, you'll supply some basic information to help the installer correctly configure your system, install the correct internationalization and location software, and set up an initial user account. The next two sections discuss your install options from the Server and Alternate Install CDs, which share the same quasi-graphical installer discussed in the introduction to this chapter, and walk you through the installation process.

NOTE **Throughout the rest of this chapter, I'll simply refer to the Ubuntu Alternate CD and the Kubuntu Alternate CD as the "Alternate CD" — the installer and the questions that it asks are the same.**

Install Options on the Server Install CD

As you might guess, the Ubuntu Server Install CD is intended for use in installing Ubuntu on a system that will be used as a server. Because many server systems are rack-mounted systems that use a system console rather than a graphical monitor, Ubuntu Server installations do not include a graphical user interface such as an X Window system window manager or desktop such as GNOME. (Therefore, there is no such thing as the Kubuntu Server CD because KDE wouldn't be installed by it anyway.) Not installing these graphics-related packages by default reduces the amount of disk space required for a basic installation and reduces the amount of software that you have to keep up-to-date.

TIP **You can always add a graphical interface to server systems later if you want to have a graphical interface available on your server system to support specific tools or for your general convenience. For example, you can add the `xubuntu-desktop` package to add the Xfce desktop system, the `ubuntu-desktop` package to add the complete GNOME desktop, the `kubuntu-desktop` package to add the complete KDE desktop, or add the `xserver-xorg` package, the window manager of your choice, and specific graphical tools. Adding software packages is discussed in detail in Chapter 23, "Adding, Removing, and Updating Software."**

Figure 3-1 shows the initial screen that is displayed when you boot your system from an Ubuntu Server Install CD, after selecting your language from the menus that are initially displayed.

The options on the Server Install CD's boot menu are as follows:

- **Install to the hard disk:** Installs a basic Ubuntu Server system to your hard drive(s) using Ubuntu's quasi-graphical installer. Installing an Ubuntu server using this installer is discussed in detail in the next section, "Installing an Ubuntu Server."
- **Check CD for defects:** Tests the integrity of the CD that is currently inserted in your system and verifies its contents.
- **Test memory:** Performs several memory tests to verify your system's memory and associated cache. You can press the letter *C* to display a configuration menu that enables you to run or rerun specific tests. Once you are finished testing your system's memory, press Escape to reboot your system and return to the Ubuntu Server Install CD boot menu.
- **Boot from first hard disk:** Bypasses booting from the Ubuntu Server Install CD and boots from your system's hard drive. This is the option to use if you accidentally left the Ubuntu Server Install CD in the CD drive of a system on which you did not intend to install Ubuntu or perform the available system tests.
- **Rescue a broken system:** Prompts for a few pieces of information such as the language that you want messages to be displayed in and your geographic location, and then boots a diskless version of Ubuntu Linux that runs from the CD and which you can use to repair corrupted disks, correct or reinstall your GRUB boot loader's configuration, and so on.

FIGURE 3-1

The Boot menu on the Ubuntu Server Install CD



TIP **If you have problems entering rescue mode, booting from the Ubuntu Desktop DVD included with this book provides the same capabilities as a rescue mode installation, but with a whizzy graphical interface.**

Selecting the default Install to hard boot menu option will install an Ubuntu Linux Server system on your hard drive(s). The next section walks you through the process of installing an Ubuntu Server after selecting the Install to hard disk option.

Installing an Ubuntu Server

Depending on the speed of your system and Internet connection, installing Ubuntu can take a little while. However, the wait is worth it. Insert the Ubuntu install CD in your CD drive and boot from it. Now do the following:

1. After booting from the CD and selecting your language from an initial set of menus, the screen shown in Figure 3-1 displays. The Server Install CD actually has three installation options, which you can display by pressing F4, which displays the pop-up menu shown in Figure 3-2. These options are as follows:
 - **Normal:** Install a standard Ubuntu Server system.
 - **Install a minimal system:** Install an Ubuntu Server system with the minimum number of packages required to boot and run. You should select this option if you want to configure your server manually, installing only the specific packages (and their dependencies) that you are interested in.
 - **Install a minimal virtual machine:** Install an Ubuntu Server system using the Ubuntu JEOS (Just Enough Operating System) configuration. (See the section of Chapter 1 entitled "Server

Users" or `www.ubuntu.com/products/whatisubuntu/serveredition/jeos` for more information on JEOS.) Like the Install a minimal system option, this option installs an Ubuntu Server system with a minimum number of packages, but further reduces the list of installed packages and installed hardware drivers to create a truly minimal system that is suitable for use as a virtual machine installation.

To proceed, select any of these options or simply press Return to leave the Normal option selected. The screen shown in Figure 3-1 re-displays. Make sure that the Install to the hard disk entry is still selected, and press Return to continue with the installation process.

FIGURE 3-2

Available configurations on the Ubuntu Server Install CD



2. The screen shown in Figure 3-3 displays, prompting you to select the language in which you would like system messages and dialogs to be displayed. You can either press Return to accept the default value of English, or use the arrow keys to scroll up and down in the list of supported languages until you find your native language, and then press Return to continue.
3. The screen shown in Figure 3-4 prompts you to specify the location where you are using your computer. This value is used to further customize the language that you specified in the previous screen by using the appropriate date and time formats, numeric conventions, and currency values for your geographic location.

FIGURE 3-3

Specifying your language

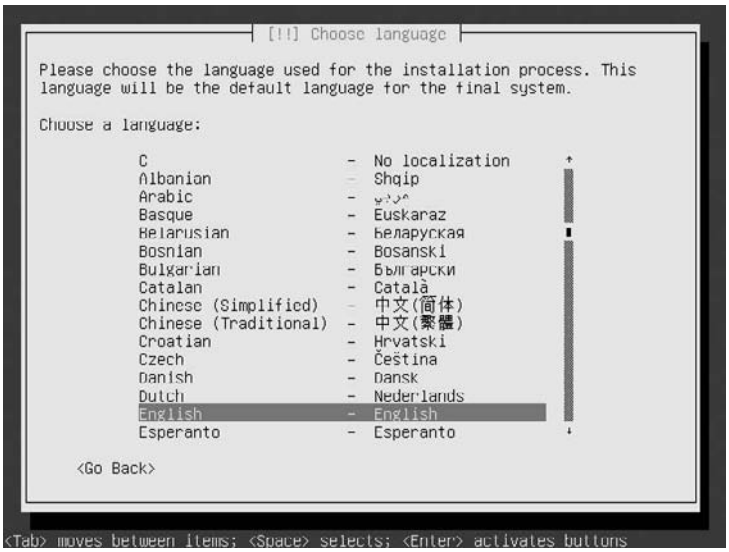


FIGURE 3-4

Specifying your geographic location



4. The screen shown in Figure 3-5 prompts you as to whether you want to try to automatically detect or manually specify the type of keyboard that you are using:
 - Selecting Yes and pressing Return displays a few additional screens that prompt you to press specific keys on your keyboard and ask if certain keys exist on your system.
 - Pressing Tab, selecting No, and pressing Return displays two dialogs that contain lists from which you can select the type of keyboard that you are using.

 If you requested auto-detection, a confirmation dialog displays after pressing the requested keys to ensure that the correct keyboard was detected. If this is correct, press Return to proceed; otherwise press Tab to select your keyboard from the lists in two successive dialogs.

FIGURE 3-5

Specifying your keyboard type



At this point, a few screens display as the Ubuntu installer looks for your CD-ROM drive, probes your system's hardware, and identifies and configures your network interface. These screens are not interactive.

TIP **If the installer has problems automatically configuring your network interface and you see an error dialog at this point, see the sidebar "Resolving Installation Problems" later in this chapter for more information.**

5. The screen shown in Figure 3-6 prompts you to enter a name for this machine. This is a one-word entry (without any domain name) that can be used to uniquely identify this system on your local network. If I am dual-booting, I tend to use the same name for the system whenever possible, to avoid confusing myself. If you want, you can accept the default name *ubuntu*, or you can enter a name for your machine (as I have in Figure 3-6), and then press Return to continue.

FIGURE 3-6

Specifying a name for your Ubuntu Server

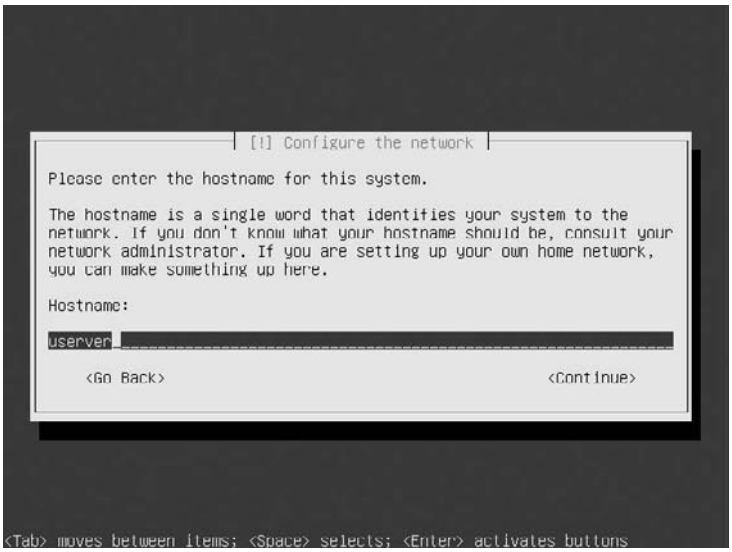


FIGURE 3-7

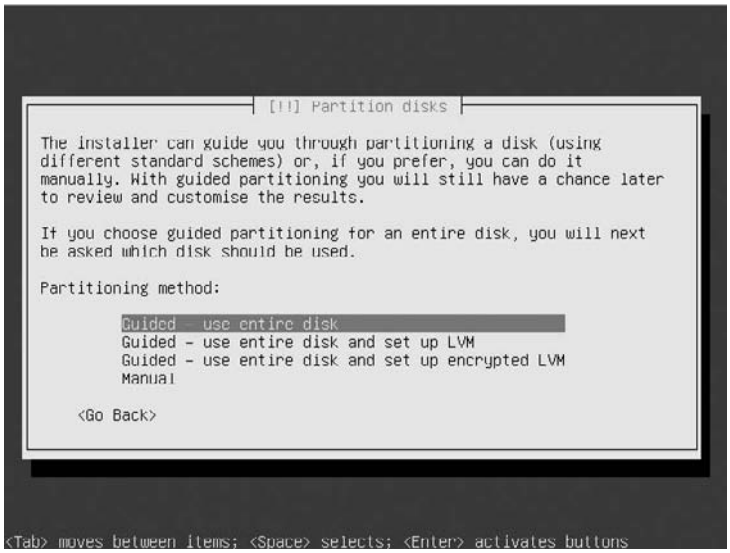
The time zone screen



6. The screen shown in Figure 3-7 enables you to specify the default time zone in which you will be using your system. Even if you're installing on a laptop, you should enter your home time zone as a base reference — you can always change this later if you travel for an extended period of time. Use the arrow keys to select the appropriate time zone for your home location and press Return to continue.
7. The installer displays several informative screens as it probes and identifies your system's hard disks. The partitioning screen shown in Figure 3-8 displays once these tests have completed.

FIGURE 3-8

The standard Ubuntu Server Partitioner screen



8. The partitioning screen shown in Figure 3-8 is the only potentially complex portion of the Ubuntu installation process. If you are installing the Ubuntu Server on a system with a single disk and you want to devote that entire disk to Ubuntu, the installation screen should look much like the screen in Figure 3-8 except for differences in device names and sizes. In this case, you have a few options:
 - Erase the entire disk and use traditional physical partitions.
 - Erase the entire disk and use logical volumes.
 - Erase the entire disk and use encrypted logical volumes.
 - Manually specify the partition layout for your system.

 If you want to manually partition your system, see the section later in this chapter entitled "Manually Specifying Your Partition Layout," and then return here once you've finished partitioning. Otherwise, select one of the other options. The advantage to using logical volumes is that if you run out of space on your existing system, you can always add another disk drive and

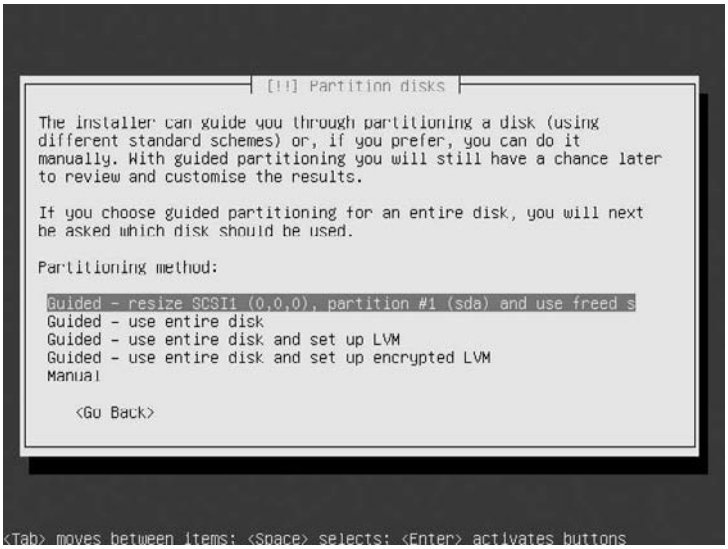
add the space that it provides to your existing logical volume. The advantage to using encrypted logical volumes is that your data will be secure if someone steals your machine and sells it on eBay (particularly common for laptops). There are plenty of ways to do these same sorts of thing using physical volumes, so unless you're sure that you want to use logical volumes, just press Return to accept the default value and proceed to the next step.

TIP

If you are installing Ubuntu on a dual-boot x86 or 64-bit PC system, you will see something more like the partitioning screen in Figure 3-9. In this case, you can select the Resize option to let the installer automatically resize your disk, or you can select the Use the largest continuous free space option to install Ubuntu on the largest unused portion of your disk if you manually resized your partitions or simply have unallocated space on the disk. If you select the Resize option, a subsequent screen displays on which you must specify the size to which you want to reduce your Windows partition. This size can be expressed as an absolute size (for example, 10GB), or as a percentage of the disk (for example, 25%).

FIGURE 3-9

The Ubuntu Partitioner on a dual-boot system

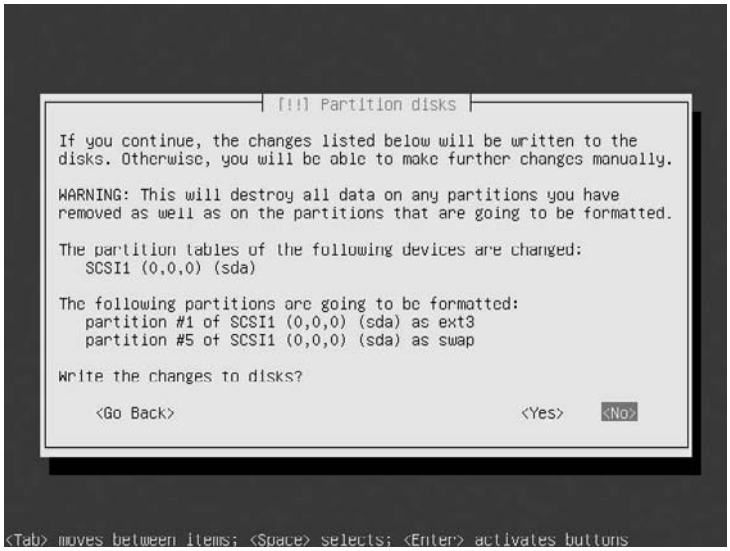


9. The Partitioner displays an additional screen that requests confirmation of the disk that you want to partition. If your system contains multiple disks, make sure that the right one is selected. Press Return to proceed. If the disk that you have selected for partitioning already contains partitions or logical volumes, an additional confirmation screen displays that gives you a final chance to avoid formatting the wrong disk by requiring confirmation before proceeding. If the correct disk is selected and you do indeed want to remove any existing data that it contains, press Return to proceed.

10. The Ubuntu Partitioner's summary screen displays, as shown in Figure 3-10. As the screen says, this operation cannot be undone. Double-check the displayed values. To accept the values that it displays, press the Tab key to select the Yes value on the screen, and press Return to continue. To make changes to the suggested partitioning scheme, make sure that you select the No value and press Return. You will be returned to the manual partitioning screens discussed later in this chapter in the section "Manually Specifying Your Partition Layout." In that case, please see that section for more information, and then return to this point of the installation process.

FIGURE 3-10

The partitioning summary screen



11. The screen shown in Figure 3-11 displays, prompting you for the name of a user to create on your system. As discussed in Chapter 4, "Basic Linux System Concepts," Ubuntu uses a special permissions scheme to give a single user (by default) the ability to perform administrative tasks on your machine. The name of the user that you enter here will have those administrative abilities, so this should almost always be your name or the name of a generic system administration user. Enter an appropriate full name and press Return to proceed.

12. The screen shown in Figure 3-12 displays. You must enter the login name for the user whose name you specified in the previous step. This screen initially displays a suggested login name based on the full name that you specified in the previous step. If you want to change this, use the Backspace key to delete the default suggestion and enter the login name that you want to use. To proceed with the default or with a custom username, press Return to continue.

FIGURE 3-11

Creating the default user account

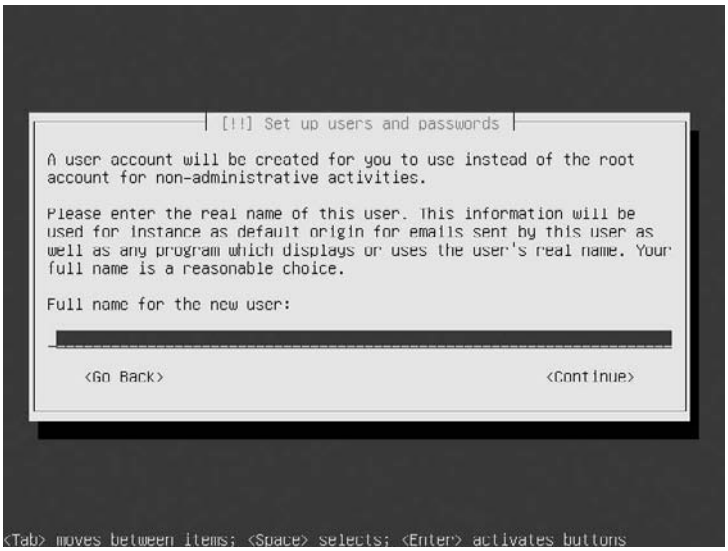
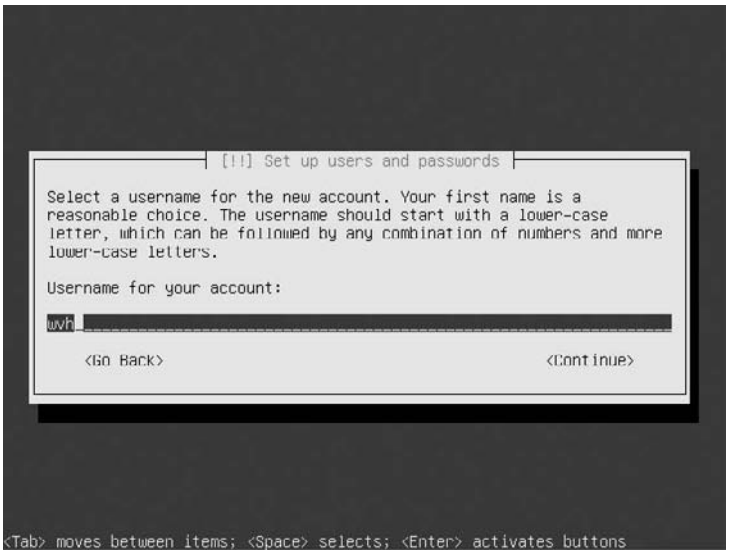


FIGURE 3-12

Creating the default login name



13. The screen shown in Figure 3-13 displays. Enter the password for the user that you are creating, following good password rules such as using a password that contains a mixture of uppercase and lowercase letters and one or more numbers, and *not* using the name or birthday of your spouse, children, or yourself. Press Return to continue.
14. After entering your password for the first time, a second password confirmation screen displays. Enter the same password that you entered on the previous screen, and press Return to continue.
15. The dialog shown in Figure 3-15 displays, asking if you want to create an encrypted private volume in which you can store documents that you do not want other users to be able to access. This volume will be created in a file, and will require a special password in order to mount that file and access its contents.

 To create this secure file volume, press Return. The dialog shown in Figure 3-16 displays. To skip creating a private volume, use the Tab key to select No and press Return. The dialog shown in Figure 3-17 displays, and you can proceed to the next step of the installation process.

 If you are creating a private volume, enter a secure passphrase in the dialog shown in Figure 3-16. This passphrase will be used to give you access to this volume in the case of errors that require manually mounting it. Press Return to proceed. A confirmation dialog displays, asking you to re-enter the passphrase to ensure that you didn't make any typos when entering it the first time. Re-enter your passphrase and press Return to proceed.
16. The dialog shown in Figure 3-17 prompts you as to the location of any HTTP proxy that you need to use in order to access the Internet to download packages. If you need to use an HTTP proxy, enter its URL in this dialog. If you do not need to use an HTTP proxy, leave this field empty. Press Return to proceed.

FIGURE 3-13

Specifying a password



FIGURE 3-14

Confirming a password



FIGURE 3-15

Optionally creating a private, encrypted volume

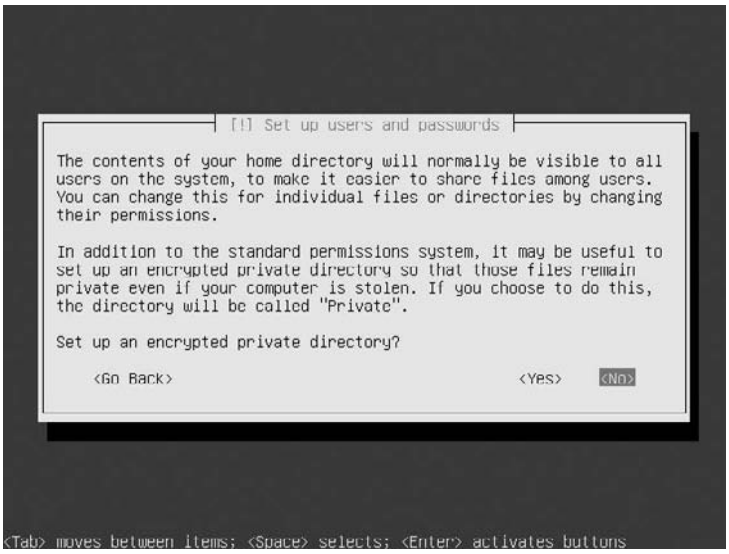


FIGURE 3-16

Specifying a passphrase for a private, encrypted volume

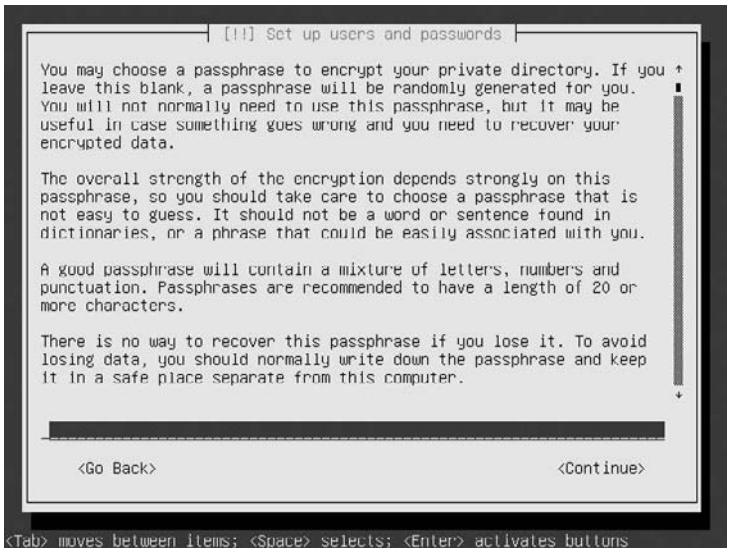


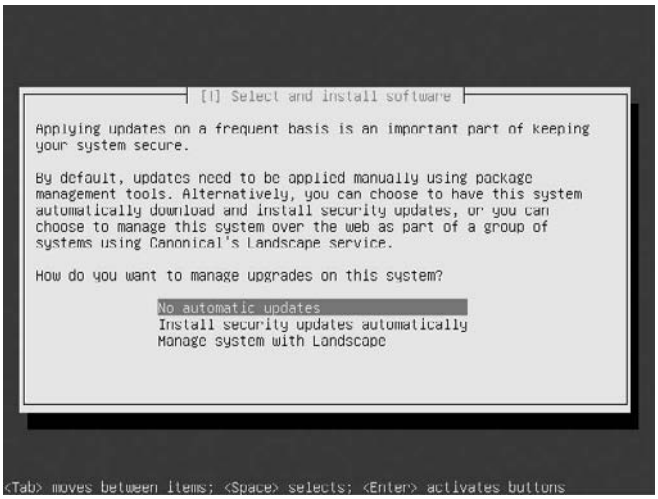
FIGURE 3-17

Identifying an optional HTTP proxy



FIGURE 3-18

Specifying an update mechanism



17. The dialog shown in Figure 3-18 displays, which enables you to specify whether the system should not download updates automatically, should download updates automatically, or should use Landscape to manage the system. (Landscape is a Web-based system management and monitoring system that is available from Canonical. For more information about Landscape, see `www.canonical.com/projects/landscape`.)

Make your selection on this screen and press Return to proceed.

18. The dialog shown in Figure 3-19 displays, which enables you to select a variety of server-related packages for installation along with the base system.

Select any packages that you want to install as part of your initial installation by using the arrow keys to select their names and pressing the spacebar to select the highlighted package. When done, press Tab to select Continue, and press Return to proceed.

19. At this point, the installer begins copying files from the CD to the hard disk, installing the default packages and any that you selected from the dialog shown in Figure 3-19. If you selected any additional packages in Figure 3-19, the installation process will display any dialogs that are necessary for their configuration. See the bonus chapters on the CD included with this book for information about the configuration requirements of many of these server packages.

TIP **If you see an error message about the Ubuntu repositories during the installation process, see the sidebar "Resolving Installation Problems" later in this section for more information. This is a nonfatal error, but is something that you'll want to correct.**

20. If you are installing an Ubuntu Server on a dual-boot system, the screen shown in Figure 3-20 displays during the package installation process. The default location to which GRUB, the Linux boot loader used by Ubuntu, is installed is to the master boot record on the primary hard drive. To install GRUB to this location, press Return.

FIGURE 3-19

Selecting server software for installation

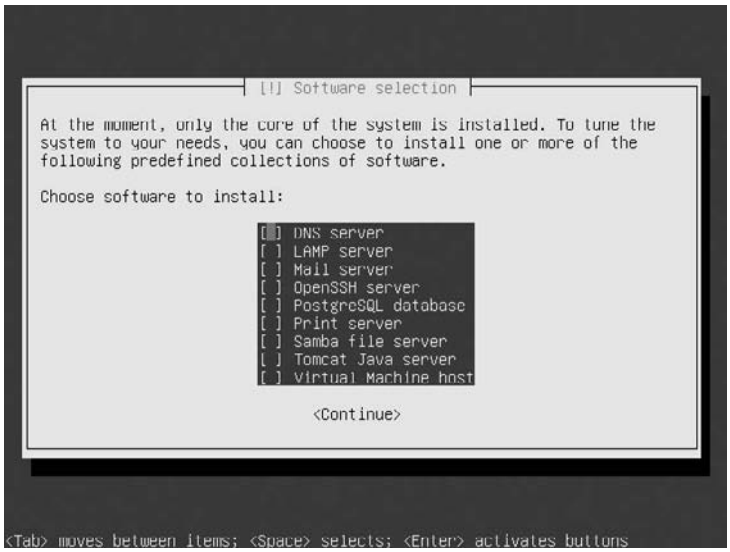
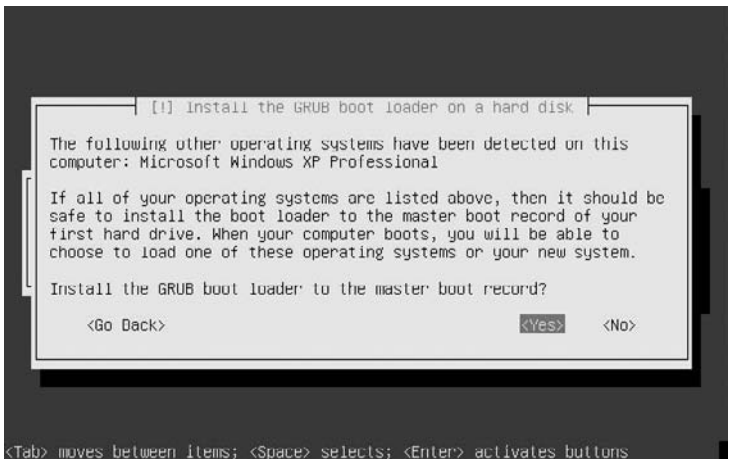


FIGURE 3-20

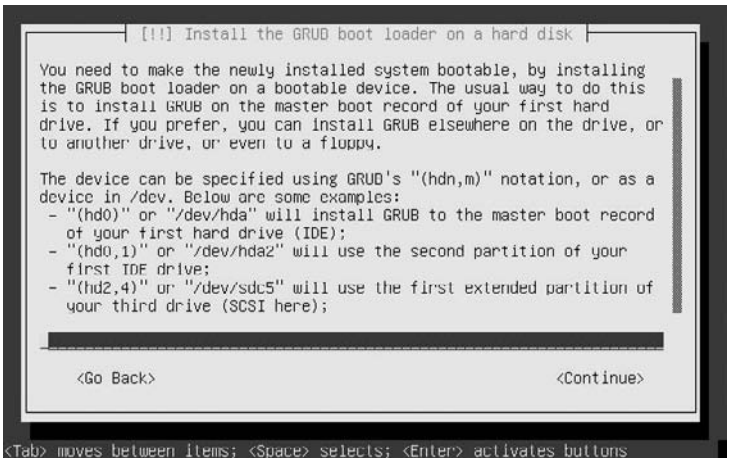
The GRUB boot loader location screen for dual-boot systems



To install GRUB to a different location, use the Tab key to select No and press Return. The screen shown in Figure 3-21 displays, in which you can specify where you want to install GRUB. Enter the value where you want to install GRUB and press Return to continue.

FIGURE 3-21

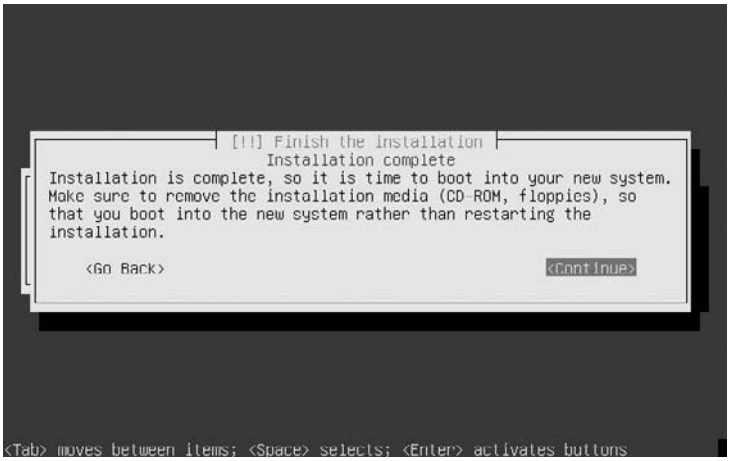
The GRUB boot loader installation screen



21. After some additional package configuration, the screen shown in Figure 3-22 displays.

FIGURE 3-22

The final installer screen



Once you see the screen in Figure 3-22, remove the install CD and press Return to reboot your computer into your newly installed Ubuntu Linux server system! Skip ahead to the section "Booting Your Server for the First Time," later in this chapter, for more information.

Resolving Installation Problems

Most of the problems that you encounter when installing any operating system, including Ubuntu Linux, are fatal. Problems encountered when reading from the installation media; accessing, partitioning, and formatting your system's hard drives; and other hardware-related problems must be resolved before you can successfully install Ubuntu Linux. However, you may encounter two fairly common problems when installing Ubuntu Linux that can easily be corrected. The first is a problem with automatic network configuration; the second is a problem accessing Ubuntu's online repositories and correctly configuring the Ubuntu software update system.

If the Ubuntu installer cannot automatically configure the network, you will see a screen stating that automatic network configuration failed and asking you if you want to manually configure the network, as shown in the following figure.



The Network autoconfiguration failed screen

Select Continue or press Return to display the screen shown in the next figure.

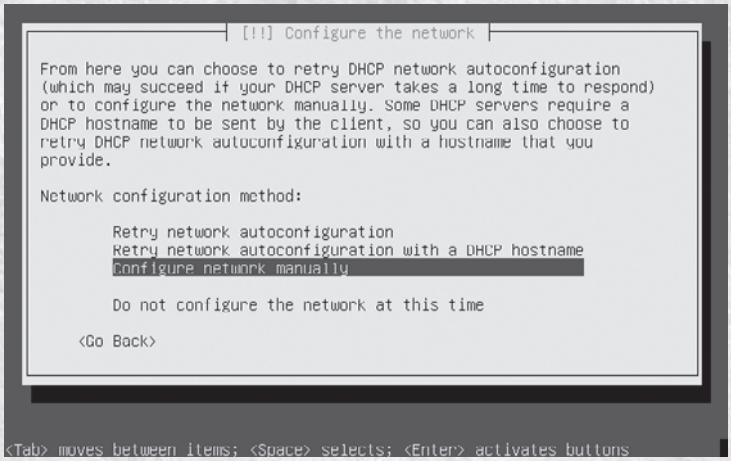
This screen enables you to retry automatic network configuration via DHCP broadcasts, retry automatic network configuration via DHCP by querying a specific host, configure your network manually, or skip network configuration altogether. Before doing anything else, check your network cabling to make sure that the machine you are installing Ubuntu on is correctly attached to the network, especially if your network is actually running a DHCP server.

You can select any of the available options, but I suggest that you select the Configure network manually option to provide your Ubuntu Server with a static IP address and related networking information. In my experience, few servers have dynamic IP addresses because this makes it more challenging to contact them remotely if a problem occurs, because you aren't guaranteed to know their IP address.

If you select the Configure network manually option, the installer displays four additional network configuration screens that respectively prompt you for the server's IP address, the netmask for that IP address, the IP address of your network gateway, and the name server(s) on your network. (If you want to enter multiple name servers, separate them by a space, not a comma.) Once these screens are displayed, the system's network interface is configured, and you are returned to the installer.

continued

continued



The manual network configuration screen

Another common problem you may encounter when installing Ubuntu Linux is in contacting and verifying all of the default Ubuntu Linux repositories, which are sites on the Internet that provide new and updated software packages for different Ubuntu releases. If problems occur when contacting the Ubuntu repositories, a dialog about this failure may display during the installation process, and certain repositories may be disabled temporarily.

If you see a message along these lines, you shouldn't be concerned. This is usually a transient, network connectivity or load problem. After you reboot, you will need to uncomment any repositories that were identified in this dialog. (Most of these problems occur when contacting the security repositories.) To uncomment these, use your favorite text editor via a command such as `sudo vi /etc/apt/sources.list`. Remove the hash mark at the beginning of each of the lines associated with the specified repositories, save the modified file, and execute the command `sudo apt-get update` to verify that the new repositories can be contacted successfully. See Chapter 23, "Adding, Removing, and Updating Software," for more information about the Ubuntu repositories and software package management utilities.

Manually Specifying Your Partition Layout

Manually specifying the layout of the partitions on your system can be useful, but is also time-consuming and requires some understanding of the Linux filesystem. Manually partitioning your system enables you to put specific directories from the Linux filesystem on their own partitions, which can improve performance, minimize the size of backups for your system (because backups are usually done on a per-filesystem or per-directory basis), and simplify future upgrades and system reconfiguration.

NOTE **By default, the Ubuntu installer automatically creates a reasonable partitioning scheme for you, which means that it creates one large partition for user and system files and another, smaller, partition to use as swap space to support virtual memory on your Ubuntu system. This section describes an alternative to the quasi-graphical Ubuntu installer's automatic partitioning scheme. If you're reading this book chapter-to-chapter, you can skip over this section if you've already installed your system.**

If you are installing Ubuntu for the first time or are relatively unfamiliar with Linux, I'd suggest letting the Ubuntu installer automatically partition your disk for you. Automatic partitioning will get you up and running much more quickly, and doesn't require that you understand the content and use of various directories on your Linux system. As you become more familiar with Ubuntu and Linux, you can always subsequently back up your user account to another machine, and then reinstall Ubuntu using manual partitioning.

The following are common directories in the Linux filesystem that are often put into their own partitions:

- `/`: The top level of the Linux filesystem. A partition must always be available through the `/` directory in order for a Linux system to boot.
- `/boot`: The directory that holds the Linux kernel and other associated system files used during the boot process.
- `/home`: The directory where user files and accounts are located.
- `/opt`: A directory where optional programs and related files are typically installed.
- `/tmp`: A directory used to hold temporary files created by user and system processes.
- `/usr`: A directory that holds applications, system files, and libraries used by the standard Linux system. On Ubuntu systems, the `/usr` directory requires approximately 1.25GB of disk space by default.
- `/usr/local`: A directory that holds applications, system files, and libraries used by a particular Linux system, but which may not be present on all Linux systems.

Most Linux systems that do not use a single partition to hold the entire filesystem create separate partitions for `/`, `/boot`, and `/home`. For more information about these directories and the structure of the Linux filesystem in general, see Chapter 4, "Basic Linux System Concepts."

Linux partitions can be grouped into two general types: swap partitions, which are used internally by the system to support virtual memory; and data partitions, in which files and other directories are located. A Linux system must have at least one area to swap to in order to function correctly — this is usually a dedicated partition, although you can also swap to a file in the filesystem if you are desperate. (See Chapter 4 for more information about partitions and virtual memory.) In this section, you will create a few basic data partitions and a single swap partition to illustrate manual partitioning — how you actually decide to partition your system is up to you.

Manually partitioning your disk involves steps like the following:

1. In the screen shown in Figure 3-14, select the Manual option. The screen shown in Figure 3-23 displays.
2. Use the arrow keys to select the entry representing the disk drive that you want to partition and press Return to continue. The screen shown in Figure 3-24 displays.
3. Use the Tab key to select Yes and press Return to continue. Figure 3-23 redisplays.
4. Use the arrow key to select the pool of available space on that disk and press Return to continue. The screen shown in Figure 3-25 displays.

FIGURE 3-23

The screen for manually partitioning disks

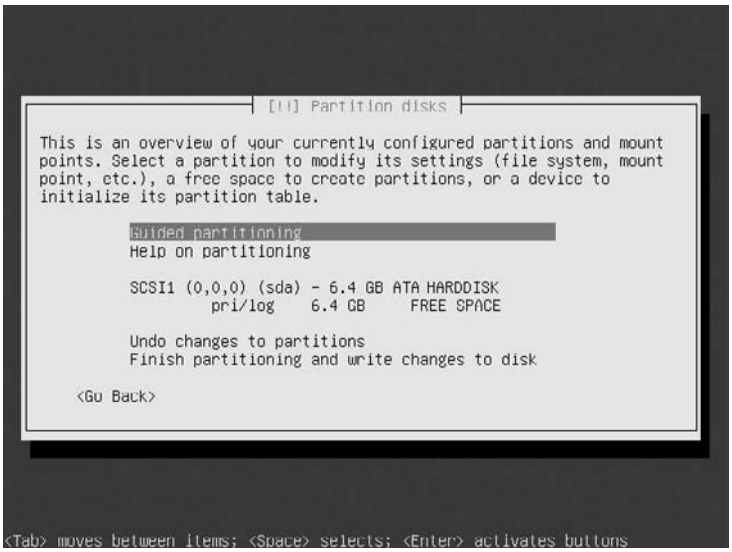


FIGURE 3-24

A warning screen regarding manual partitioning

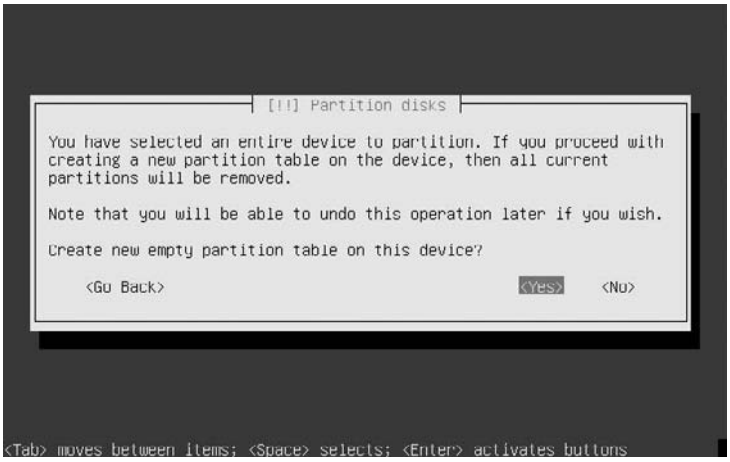
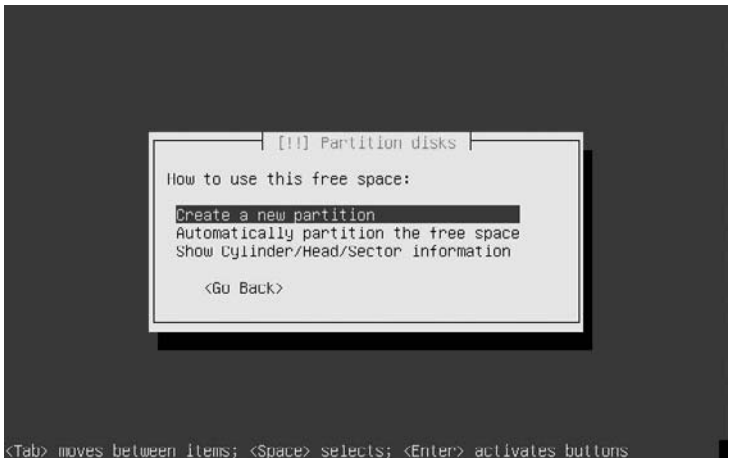


FIGURE 3-25

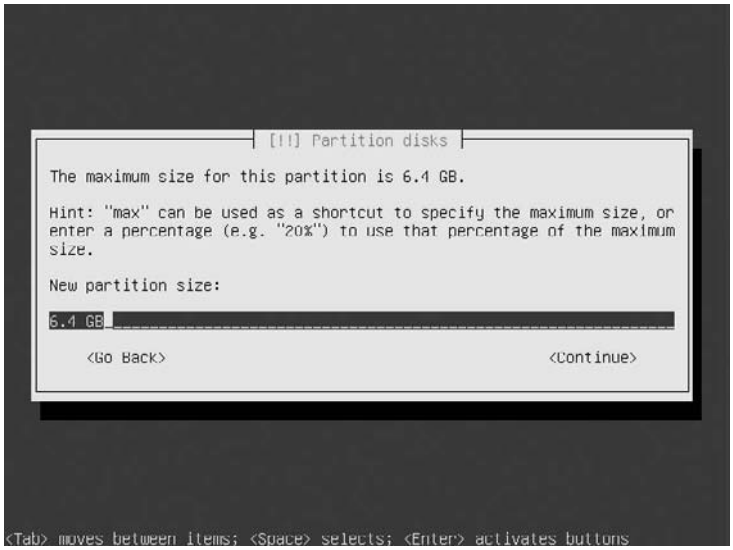
Beginning to define a new partition



5. Select the Create a new partition entry and press Return to continue. The screen shown in Figure 3-26 displays.

FIGURE 3-26

Specifying the size of your new partition



6. Enter the size that you want your new partition to have. The amount of unallocated disk space is displayed as a default value. You can specify the size of your new partition as an absolute value, such as 100MB; specify it as a percentage of the available space, such as 10%; or use the `max` keyword to create a partition of the maximum size available based on existing free space. After entering this value, press Return to continue. The screen shown in Figure 3-27 displays.

FIGURE 3-27

Specifying the type of your new partition



7. Select the type of partition that you want to create. Disks can contain up to four physical partitions — any partitions beyond that number must be logical partitions. Logical partitions can contain other partitions, while physical partitions are just that — physical portions of your disk drive that cannot be further subdivided. Unless you plan to create more than four partitions, use the arrow keys to ensure that the physical entry is selected, and press Return to continue. The screen shown in Figure 3-28 displays.
8. Unless you have a specific reason to position the new partition on a special portion of your disk, accept the default value Beginning and press Return to continue. The screen shown in Figure 3-29 displays.
9. The *mount point* for a partition is the directory through which the contents of that partition are made available to your system. (See Chapter 4, "Basic Linux System Concepts," for more information about Linux directories and mount points.) A list of common partitions and associated mount points was given at the beginning of this section. The only mandatory filesystem and associated mount point on a Linux system is `/`, the root directory of the Linux filesystem. Use the arrow keys to select the Mount Point menu item, and press Return to continue. The screen shown in Figure 3-30 displays.
10. In this example, you'll mount your new partition at `/boot`, which is the directory that holds the Linux kernel and associated system files. Use the arrow keys to select the entry for `/boot`, and press Return to continue. The screen shown in Figure 3-31 displays.

FIGURE 3-28

Specifying the location of your new partition

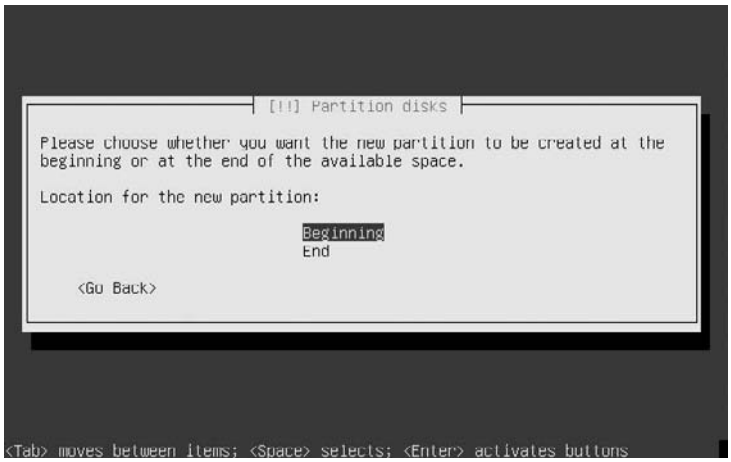


FIGURE 3-29

Specifying the type and mount point of your new partition

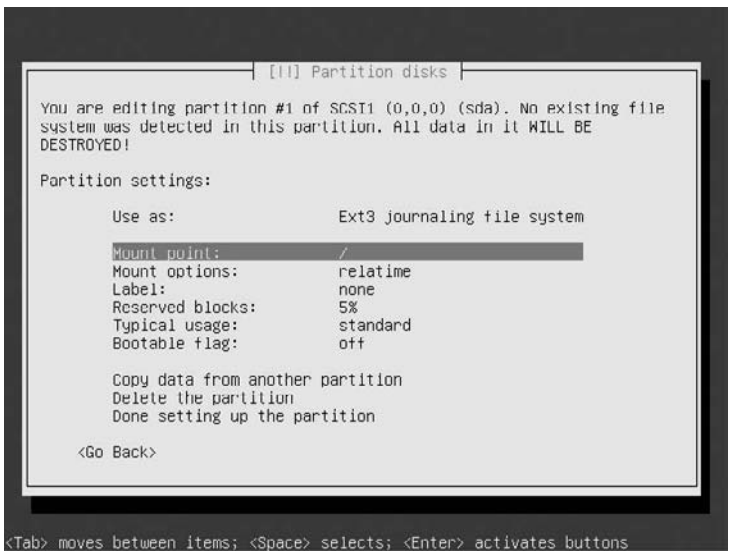


FIGURE 3-30

Specifying the mount point for your new partition

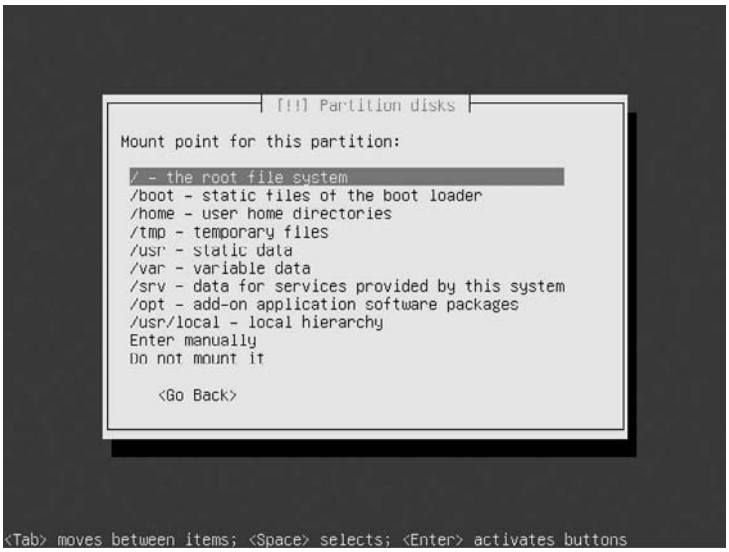
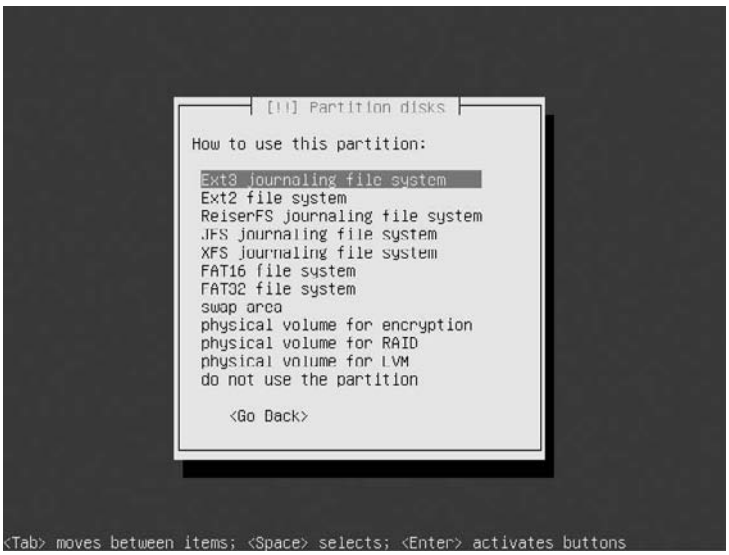


FIGURE 3-31

Specifying the partition type

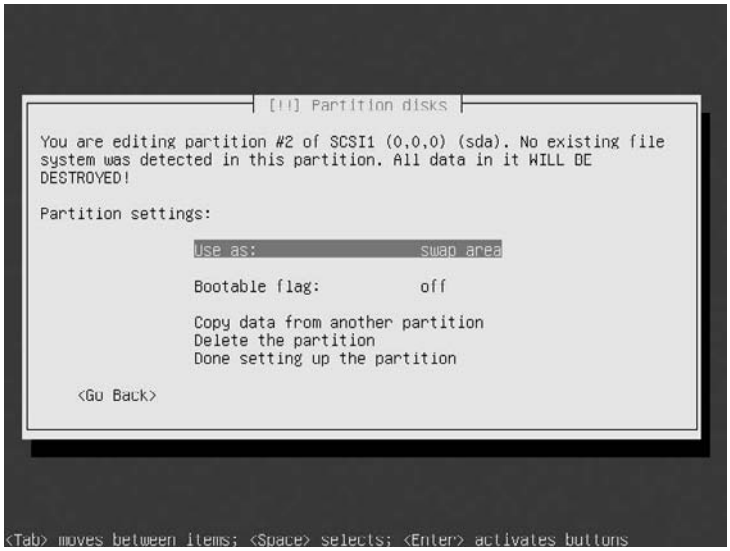


11. Linux systems support several different filesystem types. Each type of filesystem has certain characteristics that are explained in more detail in Chapter 4, "Basic Linux System Concepts." In this case, `/boot` is a data partition, and the Linux ext3 filesystem is the default (and most popular) format for data partitions. Press Return to continue. The screen shown in Figure 3-29 redisplays. Use the arrow keys to scroll down and select the Done setting up the partition option. The screen shown in Figure 3-23 redisplays, showing the new partition that you have just defined and the remaining free space for the drive.
12. Repeat Steps 3 through 11 to create additional partitions. As discussed in the previous section, filesystems that are often created on their own partitions are `/boot` (used in this example), `/`, `/home`, and a swap partition. Creating any other filesystem partitions follows exactly the same process as described previously. To create a swap partition, follow the same process as when creating other partitions, but select the swap area from the screen shown in Figure 3-33. The size of a swap partition should be approximately the same size as the amount of memory in your computer system. Figure 3-32 shows an additional screen that displays when you define a partition as a swap area. This screen merely confirms that you are creating a swap area that does not contain a filesystem that can be used to hold data files.

CAUTION **If you plan to create more than four partitions, the fourth partition that you create must be an Extended partition, in which the other partitions that you want to use can then be created.**

FIGURE 3-32

The details screen when creating a swap partition



TIP **To help guarantee that you do not encounter installation problems, you may want to ensure that the partition that you mount at / is at least 3GB in size. Following this rule should guarantee that your Ubuntu system will install correctly unless you have defined other partitions such as `/usr` with insufficient space to hold the files that are installed under that directory when it is used as a mount point.**

13. Once you have allocated all of the available free space on your system to partitions, your screen should look something like the one shown in Figure 3-33.
14. The screen shown in Figure 3-33 displays a summary of the partitions that you have defined for use on your system. Remember that there must be one partition mounted at / and one swap partition. To accept the values shown on this screen, use the arrow keys to select the "Finish partitioning and write changes to disk" option, and press Return to display the screen shown in Figure 3-34.

FIGURE 3-33

The partition summary screen

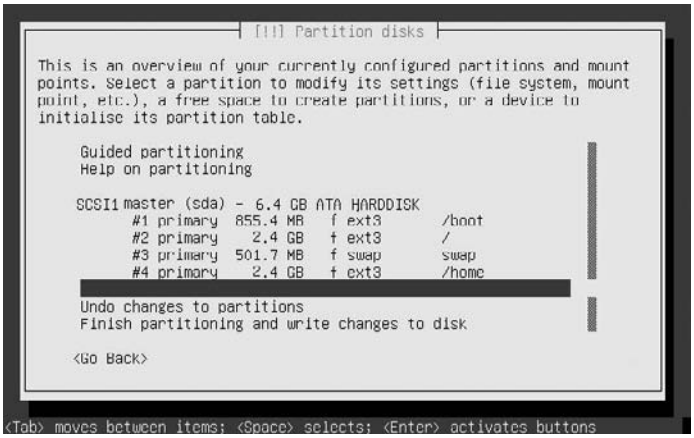
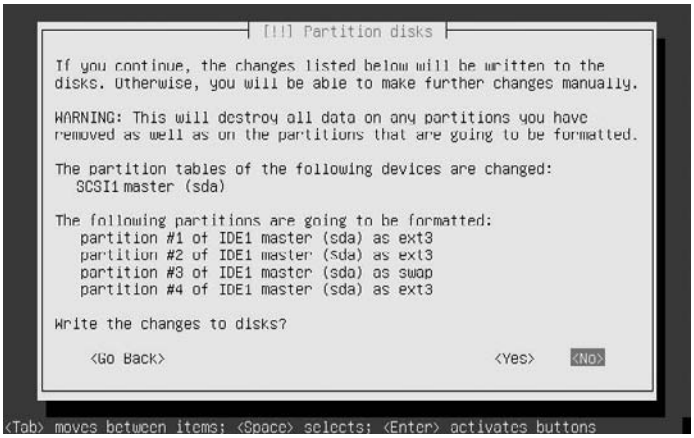


FIGURE 3-34

The final partition confirmation screen



15. Because disk partitioning cannot be undone, the screen shown in Figure 3-34 requests final confirmation of your manual partitioning selections. To accept these values and continue with the installation process, use the Tab key to select Yes, and press Return to continue.

You can now return to Step 10 of the Ubuntu Server installation process in the previous section. And congratulations — you're now at least a Linux wizard in training!

Booting Your Server for the First Time

Once the server installation process completes, the screen shown in Figure 3-22 displays. Remove the install CD and press Return to reboot your system into your newly installed Ubuntu Linux server!

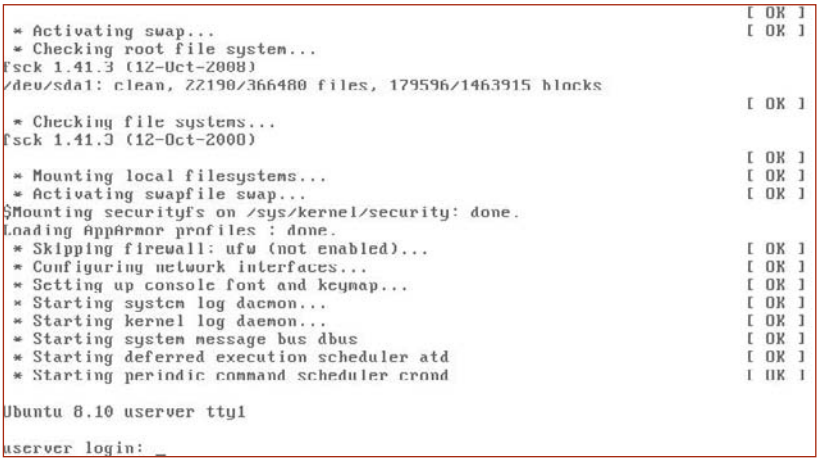
NOTE **The Ubuntu boot process differs depending on whether you have installed Ubuntu as a dual-boot system or as the only operating system on your computer. Although it is highly unlikely that you would have installed an Ubuntu Server on a dual-boot system, see the section in Chapter 2 entitled "Booting Ubuntu or Kubuntu on Dual-Boot Systems" for more information if that is the case.**

As Ubuntu boots, you will see several text messages displayed to the screen as it probes and initializes your hardware and related system software. When the boot process is complete, you will see a command-line login screen like that shown in Figure 3-35. You can now log in using the username and password that you defined during the installation process.

Congratulations — you're running an Ubuntu Linux server! For more basic information about Linux, see Chapter 4, "Basic Linux System Concepts." If you would prefer to simply start running some of the hundreds of powerful applications provided with Ubuntu Linux, check the table of contents in this book for the type of application that interests you, and skip to the related chapter.

FIGURE 3-35

The Ubuntu Server login screen



Install Options on the Alternate Install CD

Figure 3-36 shows the menu that is displayed when you boot your system from an Ubuntu Alternate Install CD.

For parity, Figure 3-37 shows the menu that is displayed when you boot your system from a Kubuntu 8.10 Alternate Install CD.

FIGURE 3-36

The boot menu on the Ubuntu Alternate Install CD



FIGURE 3-37

The boot menu on the Kubuntu Alternate Install CD



NOTE **Aside from the logo, whether they install an Ubuntu- or Kubuntu-based system, and one hidden installation option, the Ubuntu and Kubuntu Alternate CDs are identical. Throughout the rest of this section, I'll use the term Alternate CD to mean either of these CDs — as always, the choice of whether to install and run Ubuntu or Kubuntu is up to you.**

The options on an Alternate Install CD's boot menu are as follows:

- **Install Ubuntu/Kubuntu:** Installs a basic Ubuntu or Kubuntu Desktop system to your hard drive(s) using a quasi-graphical installer. Installing an Ubuntu or Kubuntu Desktop using this installer is very similar to installing an Ubuntu Server system, as discussed in detail in the previous section. The differences between using the quasi-graphical installer to install an Ubuntu Server or an Ubuntu or Kubuntu Desktop are discussed in the next section, "Installing a Desktop System in Text Mode."
- **Check CD for defects:** Tests the integrity of the CD that is currently inserted in your system and verifies its contents.
- **Test memory:** Performs several memory tests to verify your system's memory and associated cache. You can press the letter "c" to display a configuration menu that enables you to run or rerun specific tests. Once you are finished testing your system's memory, press Escape to reboot your system and return to the Alternate Install CD boot menu.
- **Boot from first hard disk:** Bypasses booting from the Alternate Install CD and boots from your system's hard drive. This is the option to use if you accidentally left an Alternate Install CD in the CD drive of a system on which you did not intend to install Ubuntu or Kubuntu, or perform the available system tests.
- **Rescue a broken system:** Prompts for a few pieces of information such as the language that you want messages to be displayed in and your geographic location, and then boots a diskless version of Linux that runs from the CD and which you can use to repair corrupted disks, correct or reinstall your GRUB boot loader's configuration, and so on.

TIP **If you have problems entering rescue mode, booting from the Ubuntu Desktop DVD included with this book provides the same capabilities as a rescue mode installation, but with a whizzy graphical interface.**

The Alternate Install CDs actually have additional installation options that you can display by pressing F4, which (on an Ubuntu Alternate Install CD) displays the pop-up menu shown in Figure 3-38. These options are as follows:

- **Normal:** Install a standard Desktop system using the quasi-graphical installer.
- **OEM Install:** Installs an Ubuntu or Kubuntu Desktop system to your hard drive(s) using Ubuntu's quasi-graphical installer. Unlike the standard installations, the OEM mode install defines an OEM user that you can use to customize the system so that you can distribute it to end users. They then run a simple command to remove the OEM user and define a standard Ubuntu or Kubuntu privileged user. Installing an OEM system is discussed in the section "Installing in OEM Mode," later in this chapter.
- **Install a command-line system:** Installs a basic Ubuntu Server system to your hard drive(s) using the quasi-graphical installer. Installing an Ubuntu Server using this installer is discussed in detail later in this chapter, in the section entitled "Installing a Command-Line system."
- **Install an LTSP server:** Only available on the Ubuntu Alternate Install CD, this option installs a Linux Terminal Server Project (LTSP) server system and a good selection of applications for use by clients of that server. Installing an LTSP server using this installer is discussed in detail later in this chapter, in the section "Installing an LTSP Server."

FIGURE 3-38

Installation options for the Ubuntu Alternate Install CD



To proceed, select any of these options or simply press Return to leave the Normal option selected. The screen shown in Figure 3-36 or 3-39 re-displays. Make sure that the Install Ubuntu (or Install Kubuntu) entry is still selected, and press Return to continue with the installation process.

The next sections explain each of these options and highlight the differences between the text-mode installers used by these options and the quasi-graphical installer described earlier in this chapter in the section "Installing an Ubuntu Server."

Installing a Desktop System in Text Mode

Selecting Install in text mode from the Ubuntu or Kubuntu Alternate Install CD's boot menu enables you to install an Ubuntu Desktop system using the text-mode installer. This menu option is useful in various scenarios, such as the following:

- You are having problems getting the appropriate Desktop CD to boot on your system because of conflicts or problems with your graphics hardware.
- You want to install the Ubuntu or Kubuntu Desktop system but use logical volumes rather than physical partitions to hold your filesystem(s).
- You want to install the Ubuntu or Kubuntu Desktop system but want to install GRUB to a location other than the master boot record of your primary hard drive.

As with the other quasi-graphical Ubuntu installers, the quasi-graphical Desktop install follows the same series of prompts and steps as discussed in the section "Installing an Ubuntu Server" earlier in this chapter. The non-interactive package installation step installs a different set of packages to your system,

including the GNOME Desktop (or the KDE Desktop from the Kubuntu Alternate Install CD), the X Window system, and several graphical tools.

The only difference between a quasi-graphical desktop install and the server install process that was discussed earlier occurs near the very end of the installation process where, before ejecting the CD and suggesting that you reboot your system, the desktop installer prompts you for information about whether you would like the system clock to be set to Universal Coordinated Time (UCT) or local time. This screen is shown in Figure 3-39.

FIGURE 3-39

Specifying system clock settings



When you see this screen, press Return to use UCT or press Tab and then Return to use local time and continue with the installation process. After this screen, the standard installation procedure continues, eventually ejecting the CD and prompting you to reboot into your new, graphical Ubuntu Desktop system.

Installing in OEM Mode

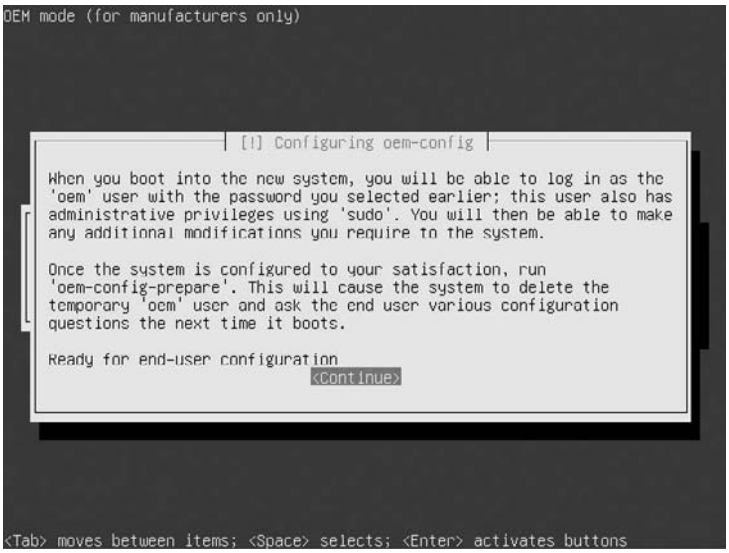
Selecting Install in OEM mode from an Alternate Install CD's Modes menu enables you to install an Ubuntu or Kubuntu Desktop system using the text-mode installer, but installs the system with a temporary privileged user (OEM). This enables OEMs to perform additional configuration on this system using this privileged user, installing custom software, installing additional packages, and removing packages as necessary, and then to redistribute the system to their customers. To turn this system into a standard Ubuntu or Kubuntu Desktop installation with a customer-specific privileged user after completing its configuration, the OEM should execute the `sudo oem-prepare` command, which will delete the OEM user and will prompt the customer to enter customer-specific user and other configuration questions the next time the system boots.

Throughout the OEM installation process, a banner in the upper-left corner of the screen identifies the fact that you are doing an OEM-mode installation. As with the other quasi-graphical installers, the quasi-graphical OEM installer follows much the same series of prompts and steps as discussed earlier in this chapter in the section "Installing a Desktop System in Text Mode." Because an OEM system is installed with a specific privileged user (OEM), the install process does not prompt you for a username and login, but only for a password for the OEM user. Like the text-mode desktop install process, the OEM mode installer prompts you for information about the resolution at which you want to display your desktop, and then continues with the installation and configuration process.

Just before the final installation summary screen, the OEM mode installer displays a screen that summarizes how OEM mode is to be used, as shown in Figure 3-40.

FIGURE 3-40

OEM mode process summary



After this screen, the standard installation procedure continues, eventually ejecting the CD and prompting you to reboot into your new, graphical Ubuntu or Kubuntu OEM mode system to complete its configuration before distribution to the OEM's customers.

Installing a Command-Line System

Selecting Install a server from the Ubuntu Alternate Install CD's boot menu enables you to install an Ubuntu Server system using the text-mode installer from the Server Install CD, and is provided on the Alternate CD for your convenience. The installation process for installing an Ubuntu Server from the Alternate Install CD is exactly the same as installing a server from the Server Install CD, as described in the section of this chapter entitled "Installing an Ubuntu Server."

Installing an LTSP Server

The Linux Terminal Server Project (LTSP, `www.ltsp.org`) is one of the most excellent Linux development efforts ever. The LTSP enables you to configure a system to act as a central server for large numbers of lightweight Linux clients, much like a Microsoft Windows Terminal Server except for the price differential. As you might expect, Windows Terminal Server and its associated per-client access licenses (CALs) cost a tremendous amount of money, while the LTSP is free for both servers and any number of clients. The LTSP is extremely popular in schools, libraries, and any other institution that needs to maximize its computing power but may have a tragically limited budget.

Selecting Install an LTSP Server from the Ubuntu Alternate Install CD's options menu enables you to install what is essentially an Ubuntu Desktop system with all of the LTSP server software built in.

NOTE **The option to install an LTSP Server is only available on the Ubuntu Alternate CD. I don't know why — seems like an oversight to me!**

For more information about the LTSP, see:

- `www.ltsp.org`: The home page for the LTSP project. Documentation for the latest LTSP releases is available at `http://wiki.ltsp.org/twiki/bin/view/Ltsp/Documentation`.
- `http://en.wikipedia.org/wiki/Linux_Terminal_Server_Project`: A good overview of the LTSP projects from Wikipedia.
- `http://k12ltsp.org/mediawiki/index.php/Main_Page`: The home page for the K12 on Linux Project, an excellent LTSP implementation that has unfortunately chosen CentOS rather than Ubuntu/Kubuntu as the foundation of its LTSP project.
- `http://ltsp.sourceforge.net`: The LTSP home page on SourceForge. There hasn't been much action there recently because the LTSP is stable and "just works." No need to fix something that isn't broken.

Summary

This chapter covered a lot of ground, explaining some basic concepts required to successfully install Ubuntu Linux as a server operating system on an existing machine or as the only operating system on a system using various specialized installation configurations via the quasi-graphical Debian installer. Regardless of the type of system that you installed or are using, by the end of this chapter, you have Ubuntu or Kubuntu Linux installed and running on your computer system.

Chapter 4 discusses basic Linux concepts that may be new to you if this is the first Linux or UNIX-like system that you have used. Chapters 5 and 6 discuss the basics of using Ubuntu with its standard graphical interface (GNOME) and window manager (Compiz). If you've opted to install Kubuntu, Chapter 7 discusses the basics of using Kubuntu 8.04 and earlier, which is based on version 3 of the KDE Desktop environment. Chapter 8 discusses the basics of using Kubuntu 8.10 or later, which is based on version 4 of the KDE desktop environment. Regardless of the Desktop environment that you select, you can look forward to a new, powerful, and completely free computer experience.

Part II

Ubuntu and Kubuntu for Desktop Users

IN THIS PART

Chapter 4
Basic Linux System Concepts

Chapter 5
Using the GNOME Desktop

Chapter 6
Using the Compiz Window Manager

Chapter 7
Using the KDE 3 Desktop

Chapter 8
Using the KDE 4 Desktop

Chapter 9
Managing E-mail and Personal Information with Evolution

Chapter 10
Managing E-mail and Personal Information with Kontact

Chapter 11
Surfing the Web with Firefox

Chapter 12
Migrating from Windows Systems

Chapter 13
Sending and Receiving Instant Messages

Chapter 14
Using Command-Line Tools

Chapter 15
Working with Text Files on Ubuntu

Chapter 16
Creating and Publishing Documents

Chapter 17
Other Office Software: Spreadsheets and Presentations

Chapter 18
Working with Graphics

Chapter 19
Working with Multimedia

Chapter 20
Would You Like to Play a Game?

Chapter 21
Consumer Electronics and Ubuntu

Chapter 4

Basic Linux System Concepts

This chapter explains basic Linux system concepts, focusing on how data is stored and organized on all Linux systems and how the Linux operating system controls access to that data and to privileged operations. Aside from specific sections on how Ubuntu deals with and grants high-level permissions, the information in this chapter applies to any Linux system and to most other UNIX-like systems.

This chapter errs on the side of caution. As impossible as it now sounds, I had never used a computer before deciding to study computer science at university. (I can be excused for this to some extent because this was a zillion years ago, just as personal computers were starting to become popular and long before DOS, Microsoft Windows, Apple Macintosh, and so on.) At any rate, I remember sitting in my first class while the professor said things such as, "Once you write your program, save it to a file" and thinking, "What the heck is he talking about?"

This chapter therefore contains some sections, such as "Working with Files and Directories" that some might find insulting in their simplicity and exploration of basic concepts that "everybody knows," which is fine with me. Everybody who knows that stuff can skip over it, but the people who just got their Ubuntu Linux boxes at the local electronics shop or from well-meaning friends or relatives can secretly read these sections, and think "Oh, I get it!" and we'll all be happier.

IN THIS CHAPTER

- Understanding files and directories
- Understanding disks and disk partitions
- Understanding logical volumes
- Linux filesystems
- Introduction to users and groups
- Privileged operations on Ubuntu systems

Working with Files and Directories

A file is nothing more than a collection of information that programs and your operating system can locate and deal with as a single unit. Files are containers for some sort of data, whether they contain a letter to your mother or parole board, a copy of one of your favorite songs (legitimate, of course), a digital photograph, or the data used by a spreadsheet to calculate the health of your

personal finances. Providing the general concept of a file as a container for related data makes it easy for applications and Linux to locate and use that information — hence the File menu that is present in most graphical applications today. This menu contains the commands that you use to open existing files, create new ones, and save changes that you have made to any file that you are working with.

Files are mandatory on a computer system — the operating system has to have a way to identify and access your data, the applications that it needs to run, any configuration files used by those applications, and so on. Similarly, when you turn it on, your computer system needs to know how to find the operating system and related configuration information that it requires to boot. However, as more and more files were required, used for different purposes, and created by different users, it didn't take long for users to look for an intermediate mechanism for organizing files, which led to the introduction of directories. Directories are simply containers for groups of related files, and they can also contain directories (or subdirectories). This is what is known as a hierarchical collection of files — the location of any specific file is described by identifying the series of directories that eventually contain the file that you're looking for. The directory that all searches start with on Linux systems is known as the *root directory* because it is the starting point for the description of how to navigate to any file that is stored locally (i.e., on your computer).

The standard analogy for files and directories is a filing cabinet — you can find any file in the filing cabinet by following a process like this: Start at the filing cabinet, open the first drawer, open the hanging file labeled Personnel, select the manila folder with your name on it, and check your healthcare enrollment form. On a computer, the process for finding the location of an analogous file might be something like this: Go to the root of the filesystem (`/`), go to the `home` directory, go to the `wvh` directory (my login name), and look for the file named `health_enrollment.txt`.

Linux systems use the "/" character to separate file and directory names, so this is often more simply expressed as, "Get the file `/home/wvh/health_enrollment.txt`." The series of directories leading to a given file is often referred to as the "path" to that file.

Linux systems that provide a graphical user interface, such as GNOME, provide graphical ways of navigating through directories, typically by clicking on them to open them and opening a window that displays graphical icons corresponding to the files that they contain. As explained in more detail in Chapter 14, "Using Command-Line Tools," any Linux command-line environment uses a command known as `cd` (which stands for change directory) to navigate through a sequence of directories. You can change to each directory in a series of directories (paths) one at a time, or you can simply `cd` to a specific target directory. In other words, the following two sequences of commands are equivalent:

```
cd /
cd home
cd wvh
```

and

```
cd /home/wvh
```

Linux systems also provide a variety of tools to list the contents of directories and provide detailed information about specific files and directories. I explain these in detail in Chapter 14. For now, I'll focus on describing how things are organized on a Linux system.

Standard Linux Directories

All Linux systems provide a standard set of core directories. The following directories are used to hold programs that must run when you are booting your Linux system, configuration files for those programs, libraries used by those programs, temporary files created by running programs, and so on:

- `/`: As the top-level directory of a Linux system, this directory must exist so that other directories can be located within it.
- `/bin`: A directory that holds core applications used by a Linux system.
- `/dev`: A directory that holds special files, known as device nodes, which are used to access any devices that are attached to your Linux system.
- `/etc`: A directory that holds system configuration information, contains the files that explain the sequence of applications that execute on a Linux system as part of its boot process, and stores configuration files for some of the applications that are executed by a Linux system.
- `/lib`: A directory that holds libraries of functions that can be called by other applications.
- `/proc`: A directory in which the Linux kernel tracks active processes and general status information.
- `/sbin`: A directory containing applications that are usually executed only by the superuser.
- `/sys`: A directory in which the Linux kernel tracks the status of system hardware and related hardware interfaces.
- `/tmp`: A directory that holds temporary files created by various applications on a running system.

You will find these standard directories on most Linux systems, regardless of the type of distribution or the size of the disk they are using.

Other Common Directories on Linux Systems

Depending on the number of files you've installed on your system and its layout, you will probably find several other directories on any Linux system. Some other commonly used directories on Linux systems are as follows:

- `/home`: A directory that holds the subdirectories where different users store their files. For example, most (if not all) of the files owned by the user "wvh" are stored in the `/home/wvh` directory (or subdirectories of that directory). The directories used by individual users to store their personal files are known as their "home directories" — this name led to user directories being stored in `/home`, not the other way around. On older UNIX and UNIX-like systems, users' home directories were stored under `/usr` (pronounced "slash user"), but are now created under `/home` to simplify system upgrades.
- `/opt`: A directory typically used when installing third-party software. This directory takes its name from the idea that it contains "optional" (i.e., non-system) software that may differ across different machines.
- `/usr`: A directory hierarchy that contains files meant to be used by normal users as they use a Linux system. The directory `/usr/bin` contains applications that users may need to execute, `/usr/lib` contains libraries used by those programs, and so on.
- `/var`: A directory that holds other directories with variable content. For example, the directory `/var/log` contains log files for system applications and events. These log files are created while a system is running, and can grow very large over time.

As you can see, a Linux system provides a large, hierarchical collection of files and directories that are organized to simplify locating certain types of files such as executables, libraries, configuration files, system status information, and so on. Now, let's look at how Linux uses various types of storage devices to store and deliver all of this information.

TIP **The directories listed in this section are fairly standard across most Linux systems, but their exact contents differ slightly from distribution to distribution. The Linux Standard Base project (LSB) was formed with the mission of standardizing directories and their contents across Linux distributions. For more information about this project, see the LSB Web site at** `www.linuxbase.org`**.**

Introduction to Linux Filesystems

The directories discussed in the previous section provide a standard way of organizing the files and directories required for a Linux system to boot, run standard applications, and so on. You'll note that the path to all of these directories begins with the / symbol. This doesn't necessarily mean that they're all located on the same hard drive or other storage device. Unlike Windows systems, which refer to different disk drives or disk partitions by using unique drive letters (`C:`, `D:`, and so on), or Mac OS X systems, which represent different disk drives or disk partitions as different volumes on the desktop, Linux systems provide what is known as a "single name space" for all of the storage devices and partitions that it can access.

This section discusses how Linux systems manage and access disks, what disk partitions are, and how disk partitions are integrated into the hierarchy of directories that you can access on a Linux system. The hierarchy of files and directories on a Linux machine is generally referred to as the Linux filesystem, although, as I'll discuss in the next few sections, the term filesystem is also used to refer to the way that data is organized on a specific disk partition. This section also introduces logical volumes, which provide a more flexible mechanism of allocating and using disk storage than simply using disk partitions, but which are most commonly only used on server systems or other systems with huge storage requirements.

Disks, Partitions, and Mount Points

The bottom line of any computer system is storing, retrieving, manipulating, and saving information. As a writer and computer systems administrator, I'm always amazed when friends who are generally computer savvy say to me, "Your laptop is so slow. Why don't you upgrade?" Frankly, if I could type faster than my primary laptop's 1.4 GHz Pentium processor could handle, I'd be posing for "Ripley's Believe It or Not!" ads or touring with the circus instead of writing books. On the other hand, if it took five minutes for me to save a chapter of whatever I'm working on or if I could never find the space to save a modified file without deleting something else, I'd drop-kick my current laptop into the river in a heartbeat and buy the newest, brightest, shiniest Linux-capable laptop available today. For what I do, being able to reliably and quickly read and write data is far more important than blazing CPU speed.

Every computer system includes some sort of storage devices such as hard disks, DVD and CD drives, CompactFlash cards, and so on, which store the operating system and your applications, and also provide some space that you can use to store the files and directories that you create. One of my favorite bits of technical writing ever is the following quote from a Hewlett-Packard manual shipped with one of their UNIX (HP-UX) workstations in the mid-1980s:

> "On a clear disk, you can seek forever."

In geek-speak, this means that without imposing some organization on the devices used to store your data, your operating system would have no idea where to look for specific files, directories, or anything

else. To store data on a disk and access that data afterward, a disk has to be prepared in a way that your operating system and associated applications can read, write, and interact with. This preparation usually consists of two steps:

1. Dividing the disk into one or more sections that can be uniquely located by the operating system. Each section of a disk is known as a partition.
2. Formatting the partitions, referred to as "creating a filesystem," in such a way that your operating system can access each partition and create files and directories there.

Disk drives are partitioned for several reasons:

- To reduce the amount of time required to locate a specific piece of data on the drive. It simply takes less time (and less location information) to find a specific piece of data in a smaller pool of information.
- To limit the amount of data that can be lost or damaged if a disk or partition becomes corrupted.
- To speed up administrative operations such as defragmentation, consistency checking, and repair (when necessary).
- To simplify administrative operations, such as backups. It's simpler to back up partitions that will fit on a single tape or other backup media because no operator intervention is required (such as switching tapes). Multiple partitions also enable you to install system files and applications on different partitions than where you store your user data. You can then back up the partition containing user data relatively frequently, without accidentally backing up a vast amount of relatively unchanging executables, system files, and so on.

As mentioned in the previous section, the term filesystem is often used to colloquially refer to the entire hierarchy of files and directories that make up a Linux system. In this section, I'm using "filesystem" in its more specific sense, which refers to the way in which a disk partition is formatted so that you can store data there and even retrieve it.

Once you format a partition and create a filesystem on it, you need to be able to make that disk partition available to your system. This is known as *mounting* that filesystem, which simply attaches that filesystem to some part of the Linux directory structure so that you can use the storage it provides. Linux systems mount filesystems on a directory so that they become a part of the standard Linux filesystem namespace, which I discuss in more detail in the section on "Mounting Filesystems," later in this chapter. Any Linux directory can act as a mount point for a filesystem, and many of the standard directories discussed in the previous section are actually mount points for separate filesystems that hold the contents of those directories. The standard Linux `/mnt` directory generally contains special-purpose subdirectories that are used as mount points for filesystems that you are planning to use temporarily.

Introducing Logical Volumes

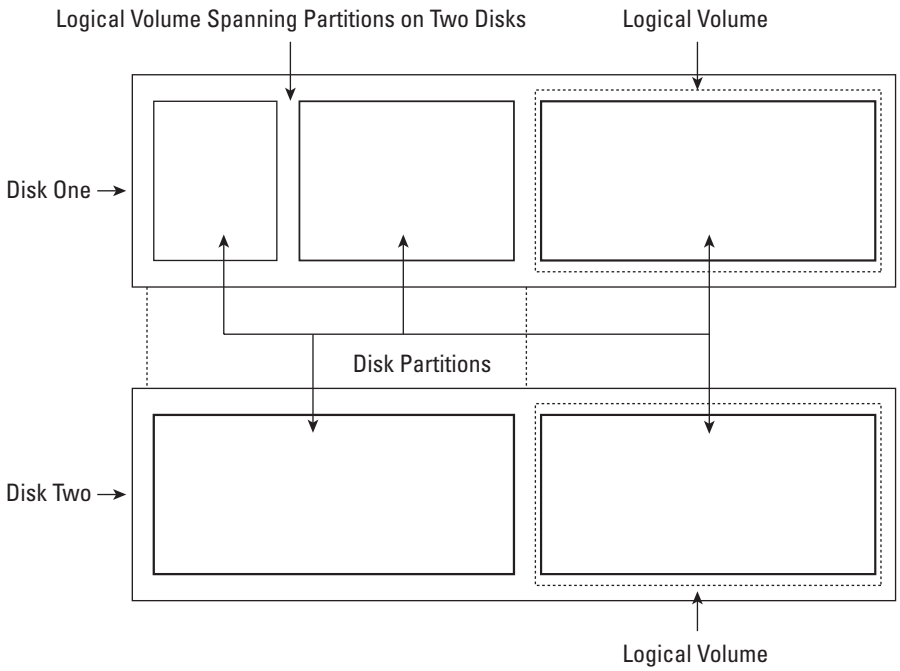
Today's disk drives are larger and faster than ever before, and both of these trends are likely to continue. However, one insurmountable problem with physical storage is that any filesystem that you create on a physical disk partition can grow no larger than the size of that partition. Although needing a partition greater than 100 gigabytes or so might be difficult to conceive of for the casual home user, placing constraints on volume size can be a serious problem for computers used in the enterprise, academic, and research communities. Identifying the largest partition that you can create at the moment is as easy as reading the label on your hard drive. Unfortunately, identifying the size of the largest partition you'll

ever need is a philosophical problem that has no answer. Although utilities are available for most modern filesystems that enable you to increase or decrease the size of a filesystem, you are still constrained by the size of the physical device on which those filesystems are located.

Logical volumes are a solution to the finite size of physical storage devices and disk partitions. A logical volume is created out of a pool of storage space, known as a volume group, that comprises partitions (known as physical volumes) from one or more hard drives, or entire hard drives themselves. You create a filesystem on a logical volume in exactly the same way that you create a filesystem on a physical partition. Figure 4-1 shows the relationships between physical storage and logical volumes.

FIGURE 4-1

Physical Disk Space and Logical Volumes



The key benefit of logical volumes appears when you run out of space on a filesystem that is contained in a logical volume. Because a logical volume is a logical entity rather than a physical one, you can easily increase its size by allocating additional storage to it from the volume group. You can then use the administrative utilities for your filesystem to increase the size of that filesystem so that it uses all of the newly added space in the logical volume. If no free space is available in the volume group, you can easily add new partitions or disks (physical volumes) to that volume's group, and then allocate some or all of that additional storage from the volume group to your logical volume. Rinse, lather, repeat.

Good examples of when you might need to increase filesystem size are simulation and image analysis projects or my research into backing up my entire CD collection onto my computer system. Without

logical volumes, you could still add new disk storage to your system by formatting new disks and partitions and mounting them at various locations in your existing filesystem, but your system would quickly become an unmanageable administrative nightmare of mountpoints and symbolic links pointing all over the place. With logical volumes, you simply add new disks to your system, add them to the appropriate volume group, and increase the size of the logical volume and the filesystem that it contains, solving the problem quickly and easily.

The standard Ubuntu and Kubuntu Desktop CDs do not support creating logical volumes during installation. Creating physical volumes, volume groups, and logical volumes on your disk(s) during installation is supported only if you are installing desktop Ubuntu or Kubuntu from their alternate desktop installation CDs or if you are installing the Ubuntu server.

For more information about creating and using physical volumes, volume groups, and logical volumes after you have initially installed your system, see Chapter 24, "Adding Hardware and Attaching Peripherals."

Local Filesystems: Standard and Journaling

Local filesystems are filesystems that are located on storage devices that are physically connected to your computer. Access to data on local filesystems is therefore fast because they are directly connected to your machine. On the other hand, local filesystems are useful only if you and any other users who need access to the data that they contain can connect to the machine on which they are located. When you need access to data that is stored in a local filesystem, connecting to the machine on which specific data is physically located generally isn't a problem in today's networked environments — you can always open an SSH (Secure Shell) or telnet connection to that machine, as long as that specific machine is up and working correctly, of course.

The most common type of filesystem used on Linux systems is the ext2 filesystem (Extended Filesystem, Version 2). I'll call this a standard filesystem because it works the way most vanilla filesystems work. When you update a file, the changes to that file are written directly to the target filesystem. If your system goes down at that exact point, your changes may or may not have been made to the target filesystem. As your system comes back up, part of its startup sequence is to check the consistency of all of its filesystems. (For information about how your system knows what filesystems it should use and check, see the section "Automatically Mounting Filesystems at Boot Time," later in this chapter.) The utility used to check the consistency of a standard ext2 filesystem is known as `fsck` (file system consistency check). The `fsck` utility traverses all of the data structures within a filesystem to make sure that all of the data in the filesystem is correctly associated with files and directories.

Running the `fsck` utility (known as fscking a filesystem) can take quite a bit of time on large partitions, which means that it takes longer for your system to become available to users. To help minimize system restart time by eliminating the need to run `fsck` (or by reducing the type of work that `fsck` has to do), a newer type of filesystem known as a *journaling* filesystem has been developed. In journaling filesystems, changes made to any file in the filesystem are initially written to a special part of the filesystem known as the journal (sometimes referred to as log) in the form of a transaction. The changes recorded in the journal are then applied to the filesystem asynchronously and transactionally. The record of these changes stays in the filesystem log until they have completed successfully, at which point the records associated with that transaction are removed.

Journaling filesystems improve system restart time because a filesystem can always be made consistent by playing any pending transactions against the filesystem. Therefore, running the `fsck` utility generally consists, at most, of executing any pending changes that are recorded in the filesystem log. Journaling filesystems can also improve overall filesystem performance because the filesystem doesn't have to wait

until a file update completes successfully, as standard (i.e., non-journaled) filesystems do. Once changes are written to the log, normal filesystem operations can continue and the filesystem can be updated asynchronously. Because the log is usually stored in a special, high-performance portion of the filesystem that uses a special format, log updates can be substantially faster than updates to the actual filesystem. Obviously, changes to any file must be completed before that file can be modified again, so it is necessary to check whether there are pending updates to a file before modifying it.

The most common journaling Linux filesystem in use today is the ext3 filesystem, which is basically the ext2 filesystem with journaling capabilities. This type of filesystem is quite popular because ext2 filesystems can be converted into ext3 filesystems with a simple change, and all filesystem repair, debugging, and analysis software written for the ext2 filesystem automatically works with the ext3 filesystem. Other popular types of journaling filesystems are ReiserFS and Reiser4 (both written by Hans Reiser and friends from namesys.com), the JFS filesystem (originally written by IBM for OS/2 and AIX, and released by IBM as an open source project, `http://jfs.sourceforge.net/`), and XFS (originally written by SGI for IRIX and released by SGI as an open source project, `http://oss.sgi.com/projects/xfs/`).

For a bit of humor, here is the man page for the `fsck` utility for the XFS journaling filesystem:

```
fsck.xfs(8)                                                fsck.xfs(8)
NAME
        fsck.xfs - do nothing, successfully

SYNOPSIS
        fsck.xfs [ ...]

DESCRIPTION
        fsck.xfs is called by the generic Linux fsck(8) program at
        startup to check and repair an XFS filesystem.  XFS  is  a
        journaling  filesystem  and  performs recovery at mount(8)
        time if necessary, so fsck.xfs simply exits  with  a  zero
        exit status.

FILES
        /etc/fstab

SEE ALSO
        fsck(8), fstab(5), xfs(5).
```

Well, that certainly cuts down on `fsck` and restart time!

Network Filesystems

Network filesystems are filesystems that are stored on remote systems (usually known as file servers) and can be mounted on your local machine just like a local filesystem. Networked filesystems provide users with the freedom to access their information from any system on which they can log in and which has access to those networked resources.

Network filesystems offer several advantages over local filesystems:

- They reduce the chance that the failure of a single machine will prevent you from accessing your data. Most networked filesystems enable you to log in on multiple machines and access your data in exactly the same way.
- They provide central locations for data that must or should be shared among all users.

- They simplify accessing existing data from faster systems. Suppose that you have written an application to test CPU and memory performance, or that your work or research depends on CPU and memory-intensive calculations. Running your application on a faster, more powerful machine is as simple as logging in on that machine and running the application from the networked filesystem.
- They provide the opportunity to centralize administrative operations such as backups.
- They promote interoperability and flexibility. You can usually access networked filesystems from systems running Linux, Microsoft Windows, Mac OS X, and so on. This makes it easy for you to use the software and hardware that is best suited to your desktop requirements and still access the same data on the networked filesystem.

A number of networked filesystems are available for Linux. The most popular of these is NFS, the Network File System originally written by Sun Microsystems and subsequently supported on almost every other UNIX-like system since the dawn of the workstation. NFS is actively supported by Apple's Mac OS X, and NFS clients (and servers) are even available for Microsoft Windows systems. The other most commonly used network filesystem is, ironically, Microsoft's SMB (Server Message Block) filesystem, which has been renamed CIFS (Common Internet File System) in an attempt to hide its proprietary roots and the fact that it had to be reverse-engineered to be accessed from Linux and other UNIX-like systems. SMB/CIFS is supported on Linux systems by a suite of client and server software known as Samba (`www.samba.org`). Samba clients enable Linux systems to mount and access SMB/CIFS filesystems exported by Microsoft Windows systems. Samba servers enable Linux systems to export their filesystems so that they can be mounted and accessed from machines running different versions of Microsoft Windows. Linux systems also support connectivity to Novell Netware filesystems through the Linux NCP — Novell Core Protocol — utilities (`www.novell.com/coolsolutions/feature/15350.html`) and pre-OS X MacOS Apple filesystems via the Linux `netatalk` and `afpfs` utilities (`http://netatalk.sourceforge.net/` or, historically, `www.anders.com/projects/netatalk/`).

Working with Partitions and Filesystems

Like Ubuntu, most Linux systems automatically identify the disks in your system and propose a reasonable partitioning scheme as part of the installation process. Similarly, when you attach a new disk to your system or add one internally, Ubuntu displays a dialog box asking if you want to configure that disk. Partitioning during the installation process is explained in Chapter 2, "Installing Ubuntu and Kubuntu;" using graphical partitioning utilities after your system is already running is explained in detail in Chapter 24, "Adding Hardware and Attaching Peripherals."

Even after you partition and mount a disk, or you attach an existing disk that has already been partitioned and formatted, your system can't access it until you mount it somewhere. This section discusses how to temporarily mount a disk using the Linux `mount` command, the most popular options that are available for that command, and how to mount a specific disk as part of your system's startup process by adding it to your system's filesystem configuration file, `/etc/fstab`.

Mounting Filesystems

Once you've attached a disk, you can always query the filesystems that it contains by using the Linux `fdisk` utility's `-l` (list) option, as in the following example:

```
$ sudo fdisk -l
```

```
Disk /dev/sda: 250.0 GB, 250059350016 bytes
255 heads, 63 sectors/track, 30401 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/sda1               1        1306    10490413+  83  Linux
/dev/sda2            1307        2090     6297480   82  Linux swap
/dev/sda3   *        2091        2122      257040   83  Linux
/dev/sda4            2123       30399   227135002+   f  W95 Ext'd (LBA)
/dev/sda5            2123        2645     4200966   83  Linux
/dev/sda6            2646        6562    31463271   83  Linux
/dev/sda7            6563       10478    31455238+  83  Linux
/dev/sda8           10479       30399   160015401   83  Linux

Disk /dev/hde: 250.0 GB, 250059350016 bytes
255 heads, 63 sectors/track, 30401 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

    Device Boot      Start         End      Blocks   Id  System
/dev/hde1               1       30401   244196001   83  Linux

Disk /dev/hdf: 250.0 GB, 250059350016 bytes
255 heads, 63 sectors/track, 30401 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/hdf1               1       30401   244196001   83  Linux
```

By default, the `fdisk` command lists all partitions on all disks that it can detect in your system. You can also use the `fdisk` command to list the partitions on a specific device, as in the following example:

```
$ sudo fdisk -l /dev/sda
Password:

Disk /dev/sda: 250.0 GB, 250059350016 bytes
255 heads, 63 sectors/track, 30401 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

    Device Boot      Start         End      Blocks   Id  System
/dev/sda1               1        1306    10490413+  83  Linux
/dev/sda2            1307        2090     6297480   82  Linux swap
/dev/sda3   *        2091        2122      257040   83  Linux
/dev/sda4            2123       30399   227135002+   f  W95 Ext'd (LBA)
/dev/sda5            2123        2645     4200966   83  Linux
/dev/sda6            2646        6562    31463271   83  Linux
/dev/sda7            6563       10478    31455238+  83  Linux
/dev/sda8           10479       30399   160015401   83  Linux
```

Note that the primary argument to this command is the base name of the disk that you want to query, not the name of a specific partition. You can get similar information, although in a slightly less usable form, by examining the file `/proc/partitions`, which a running system uses to track all disks and

partitions that are currently in use on your system. The contents of this file are the following on the same system from which the previous example was taken:

```
$ cat /proc/partitions
major minor  #blocks  name
    8     0  244198584 sda
    8     1   10490413 sda1
    8     2    6297480 sda2
    8     3     257040 sda3
    8     4          1 sda4
    8     5    4200966 sda5
    8     6   31463271 sda6
    8     7   31455238 sda7
    8     8  160015401 sda8
   33     0  244198584 hde
   33     1  244196001 hde1
   33    64  244198584 hdf
   33    65  244196001 hdf1
```

Once you've identified the partitions on your new disk (and presumably know which ones you want to use), mounting them using the `mount` command is easy. You will always need to explicitly mount filesystems that aren't listed in the file `/etc/fstab` (explained in more detail in the next section).

For example, the following command mounts the partition `/dev/sda5` on the mount point `/mnt/tmp` (which must already exist):

```
$ sudo mount /dev/sda5 /mnt/tmp
```

If you're mounting a partition that contains an ext2 filesystem or a type of filesystem that your system can auto-detect, you don't need to specify the type of filesystem that the partition/device uses. If it is something non-standard, you will have to identify the type of filesystem using the mount command's `-t` option, followed by the name of that type of filesystem, as in the following example, which mounts the XFS filesystem on `/dev/sda6` on the mount point `/mnt/xfs`:

```
$ sudo mount -t xfs /dev/sda6 /mnt/xfs
```

The `-t` option actually causes the `mount` command to first look for a special, filesystem-specific mount command, `/sbin/mount.type`, where `type` is the name of that type of filesystem. These filesystem-specific commands are usually installed along with the other administrative utilities for different types of filesystems using the Synaptic Package Manager or `apt-get`, as explained in Chapter 23, "Adding, Removing, and Updating Software."

Supported types of filesystems that you can specify using the `-t` option are `dfs`, `affs`, `autofs`, `coda`, `coherent`, `cramfs`, `devpts`, `efs`, `ext`, `ext2`, `ext3`, `hfs`, `hpfs`, `iso9660`, `jfs`, `minix`, `msdos`, `ncpfs`, `nfs`, `ntfs`, `proc`, `qnx4`, `ramfs`, `reiserfs`, `romfs`, `smbfs`, `sysv`, `tmpfs`, `udf`, `ufs`, `umsdos`, `usbfs`, `vfat`, `xenix`, `xfs`, and `xiafs`.

Some filesystems require additional options, such as NFS filesystems, SMB/CIFS filesystems, and so on. These options are highly filesystem-specific, so I won't bore you with all of them here. Instead, I'll focus on the options for a SMB/CIFS filesystem because that is probably the most common networked filesystem that you'll want to mount. Some other mount options that are commonly used in the `filesystem/mount` configuration file are discussed in the next section.

When mounting a networked SMB/CIFS filesystem, you will need to specify the name of the Windows user that you want to access the filesystem as, your Windows password, and also the identity of the Linux user who should appear to own the files in the filesystem, so that you write to the remote filesystem as needed. The following is an example of specifying these in a `mount` command:

```
//192.168.6.66/share /mnt/terastation
```

This command mounts the `smbfs` filesystem called `share` on the host `192.168.6.66` on the Linux directory `/mnt/terastation`, using the Windows username of `wvh` and whose password is `mypassword`. The mounted filesystem will appear to be owned by the user `wvh`, which means that I'll have the ability to write anywhere on that filesystem that isn't protected by Windows ACLs or some similar mechanism.

Automatically Mounting Filesystems at Boot Time

When your Ubuntu system boots, the boot block on your primary disk identifies the filesystem that contains the second-stage boot loader, which uses a configuration file to determine what kernels and associated boot options are available on your system. Older Linux systems used a boot loader called LiLo (Linux Loader), while Ubuntu and most other modern Linux systems use a boot loader called GRUB (Grand Unified Boot Loader).

During the boot process, the root filesystem is initially mounted read-only for standard processes so that its consistency can be verified. Once this is done, it is remounted in read-write mode, and your system verifies the existence and consistency of any other filesystems that it will be using. The list of filesystems that are available to your system is contained in the file `/etc/fstab` (File System Table).

Each line in the `/etc/fstab` file provides information about one of the filesystems that should be available to your system. To add a new filesystem so that it will always be mounted when your system boots, or to customize how your system interacts with a specific type of device, you need only add or customize an entry for that filesystem in the `/etc/fstab` file. A sample section of the `/etc/fstab` file on one of my systems looks like the following (this file will be different on your system):

```
$ cat /etc/fstab
LABEL=/     /              ext3        noatime,acl,user_xattr 1 1
/dev/sda3   /boot          ext3        acl,user_xattr         1 2
/dev/sda8   /home          ext3        noatime,acl,user_xattr 1 2
/dev/sda5   /tmp           ext3        acl,user_xattr         1 2
/dev/sda6   /usr           ext3        noatime,acl,user_xattr 1 2
/dev/sda7   /usr/local     ext3        noatime,acl,user_xattr 1 2
/dev/sda2   swap           swap        sw                     0 0
/dev/hde1   /opt2          ext3        noatime,acl,user_xattr 0 0
devpts      /dev/pts       devpts      mode=0620,gid=5        0 0
proc        /proc          proc        defaults               0 0
sysfs       /sys           sysfs       noauto                 0 0
/dev/hdd    /media/cdrom0  udf,iso9660 user,noauto            0 0
/dev/fd0    /media/floppy0 auto        rw,user,noauto         0 0
```

The fields in each `/etc/fstab` entry (i.e., each line) are as follows:

- The first field is the device or remote filesystem to be mounted. This is usually the Linux device file for the partition that is to be mounted, but can also be an entry of the form `hostname:directory` for networked filesystems such as NFS. `ext2` and `ext3` filesystems can

also be identified by the name that they were assigned in the filesystem volume label when the filesystem was created. For example, the entry `LABEL=/` in the example `/etc/fstab` file could be replaced with `/dev/sda1` because that is the disk partition where my root filesystem actually lives. However, using labels is more flexible than using specific partition device files because the device file associated with a specific partition may change if the disk containing that partition is moved to another system or if other disks are added to an existing system.

- The second field is the directory on which the specified filesystem should be mounted. For special types of partitions that should not be mounted, such as swap partitions, this field should contain the entry `none`.
- The third field identifies the type of filesystem. Common entries in this field are `ext2` (the standard Linux local filesystem type), `ext3` (the journaling version of `ext2`), `vfat` (a 32-bit Windows FAT partition), `iso9660` (the standard CD-ROM filesystem), `nfs` (networked filesystems using Sun's NFS protocol), and `swap` (swap space). If a filesystem is not currently used but you want to keep an entry for it in `/etc/fstab`, you can put the word `ignore` in this field, and that filesystem will not be mounted, checked for consistency, and so on.

 The types of filesystems that are compiled into your kernel are listed in the file `/proc/filesystems`, but this can be misleading. Your kernel usually also supports other types of filesystems, but as loadable kernel modules rather than being hardwired into the kernel. For example, the `/proc/filesystems` file on my system contains the following entries:

```
        ext3
nodev   sysfs
nodev   rootfs
nodev   bdev
nodev   proc
nodev   sockfs
nodev   debugfs
nodev   securityfs
nodev   pipefs
nodev   futexfs
nodev   tmpfs
nodev   inotifyfs
nodev   eventpollfs
nodev   devpts
        ext2
nodev   ramfs
nodev   hugetlbfs
        minix
        iso9660
nodev   nfs
nodev   nfs4
nodev   mqueue
nodev   rpc_pipefs
        reiserfs
nodev   usbfs
nodev   subfs
        vfat
        hfsplus
nodev   smbfs
```

- Filesystems whose types are prefaced by a `nodev` entry are not associated with physical devices, but are used internally by applications and the operating system.
- The fourth field contains a comma-separated list of any options to the `mount` command that should be used when the filesystem is mounted. Many mount options are filesystem-specific, but some common ones are the following:
 - `async`: Writes to the filesystem should be done asynchronously.
 - `auto`: The filesystem should be automatically mounted when detected or when a command such as `mount -a` is executed.
 - `defaults`: The default options `-async`, `auto`, `dev`, `exec`, `nouser`, `rw`, and `suid` are used.
 - `dev`: The character or block device containing the filesystem is local to the system.
 - `exec`: You can execute programs, scripts, or anything else whose permissions indicate that it is executable — from that filesystem.
 - `gid=value`: Set the group ID of the mounted filesystem to the specified numeric group ID when the filesystem is mounted.
 - `noauto`: Don't automatically mount when a filesystem is detected or when the command `mount -a` is issued. Usually used with removable media such as floppies and CD-ROMs.
 - `nouser`: You must be root to mount the filesystem — the filesystem can't be mounted by any non-root user.
 - `owner`: The ownership of the filesystem is set to the user who mounted it — usually root if the filesystem is automatically mounted by the system.
 - `ro`: Mount the filesystem read-only.
 - `rw`: Mount the filesystem read-write.
 - `suid`: Allow programs on the filesystem to change the user's user or group ID when they are executed if their permission bits indicate that they should do this. Be very careful when using this option with imported filesystems that you don't actually administer because running a program that sets the UID to root is a common way of hacking into a system.
 - `uid=value`: Set the user ID of the mounted filesystem to the specified numeric user ID when the filesystem is mounted.

 For more information on generic options available to the `mount` command, see the man page for the `mount` command in section 8 of the online Linux manual (by using the `man 8 mount` command).
- The fifth field is used by the `dump` command, a standard Linux/UNIX filesystem backup command, to identify filesystems that should be backed up when the `dump` command is executed. If the fifth field contains a `0` (or is missing), the dump command assumes that the filesystem associated with that `/etc/fstab` entry does not need to be backed up.
- The sixth field is used by the Linux/UNIX filesystem consistency checker (discussed in the next section) to identify filesystems whose consistency should be verified when the system is rebooted, and the order in which the consistency of those filesystems should be checked. If the sixth field contains a `0` (or is missing), the `fsck` program assumes that the filesystem associated with that `/etc/fstab` entry does not need to be checked.

The contents of this file may seem complex at first, but over time you'll get used to creating and editing them. Many of the graphical utilities provided for manipulating disk partitions automatically create these entries for you — or you can always do what I do, which is to copy an existing, but similar, entry and then modify it as needed for the partition(s) I'm adding.

Automatically Mounting Removable Media Filesystems

Beyond the entries in your `/etc/fstab` file, Ubuntu also provides special support to simplify interacting with removable media on your system. Although discussed in more detail elsewhere in this book, it is worth mentioning this here to complete this section's overview of storage device handling on Ubuntu systems.

Ubuntu's automatic recognition of removable media is done via the gnome-volume-manager application, which is configured using the gnome-volume-properties application. You can start this configuration application by selecting Removable Drives and Media from the System menu's Preferences submenu.

Once configured, the gnome-volume-manager makes it easy for your system to automatically recognize CDs, DVDs, CompactFlash cards, and other common removable media, and automatically mount them for you in appropriate mount points under your system's `/media` directory. Using and customizing the gnome-volume-manager's recognition and handling of these types of media is discussed in detail in Chapter 24, "Adding Hardware and Attaching Peripherals."

The gnome-volume-manager's support for devices that are often attached to your system while it is running extends beyond just storage devices — it includes peripherals such as printers, mice, keyboards, tablets, and so on. It also makes it easy for your system to automatically recognize consumer electronics devices such as digital cameras and personal digital assistants (PDAs) so that you can easily transfer files back and forth to the storage media that these devices use. This aspect of the gnome-volume-manager is discussed in detail in Chapter 21, "Consumer Electronics and Ubuntu."

Understanding Linux Permissions

If you're sharing your Ubuntu system with other users, it's useful to understand how Linux systems protect files and directories so that they can be accessed by only the people that you want to have access to them. Similarly, if you have installed your own Ubuntu system and are therefore responsible for taking care of it ("system administration" in geek-speak), you need to know how Ubuntu systems ensure that only specific, authorized users can perform privileged tasks. For example, it would be potentially inconvenient if random users could format disk drives, reconfigure your system's connection to networks such as the Internet, and so on. Although you'd hope that no one would do these sorts of things maliciously, it's easy enough to accidentally click OK when you are exploring system configuration applications, thinking perhaps that this would simply exit from the application.

Linux supports the traditional permission model used by all UNIX-like systems, users, and groups, with a few interesting twists that have been introduced to foster Ubuntu's goals of usability and user-friendliness. If you've used other Linux systems in the past, it's easy enough to adapt to doing things the Ubuntu way, and if you're new to Linux, you'll find the Ubuntu model for privileged commands to be quite easy to understand, configure, and use.

NOTE **Ubuntu provides an easy-to-use graphical application for creating and managing the users on your Ubuntu systems. This section focuses on the underlying concepts of Ubuntu permissions — using Ubuntu's graphical application for creating and managing Ubuntu users and groups is discussed in detail in Bonus Chapter 4, "Managing Users, Groups, Authentication, and Advanced Permissions."**

Basic Concepts: Users and Groups

All Linux systems provide two basic administrative entities that are used to determine who has access to what files and who can perform specific, privileged operations. Each person who can log in on a Linux system does so via a user account, which consists of a name and password, and which has a specific home directory that contains that user's configuration data as well as any files that they create. The file `/etc/passwd` (known as the password file, but why type a few extra letters if you don't have to) contains a list of every user who has an account on that specific machine.

All Linux users belong to one or more groups. A *group* is an administrative entity that makes it easier for multiple users to access the same sets of files. Information about the groups that are defined on your Ubuntu system and the users that belong to those groups is stored in the file `/etc/group`, which is a text file with easy-to-understand entries. As explained later in this section, file and directory permissions can be set such that users who are members of a specified group can read and/or write files in shared directories, which other users on the system still cannot access. This makes it easy to set up collaborative projects or to simply share information with selected other users on your systems.

NOTE **Most Linux systems use the local `/etc/passwd` and `/etc/group` files for authentication and group membership information, but systems that use networked authentication mechanisms can get these types of information from other systems on your network. For example, systems that use NIS, the Network Information Service that is used in many NFS-based Linux environments, can contain entries in the password and group files that tell your system to check NIS for user, group, and authentication information. Linux also supports PAMs (Pluggable Authentication Modules), which enable the system to contact networked sources of these types of information such as LDAP (Lightweight Directory Access Protocol) servers, Kerberos servers, and even Microsoft Windows authentication servers. PAMs aren't discussed here because they are very special, advanced cases, and this section is only intended as an introduction to Ubuntu and basic Linux security. NIS is discussed in Bonus Chapter 4, "Managing Users, Groups, Authentication, and Permissions."**

My entry in the `/etc/passwd` file on one of my Ubuntu systems is as follows:

```
wvh:x:1000:1000:William von Hagen,,,:/home/wvh:/bin/bash
```

The fields in this entry are as follows:

- My login name.
- The password field. In this case, an `x` means that password information is actually stored in the file `/etc/shadow`.
- My user ID, which is the numeric value associated with my login name.
- The numeric group ID of the default group of which I am a member.
- A field containing my full name and other text information such as office location and office and home phone numbers.
- My home directory.
- The application that runs when I log in, in this case a shell.

A few sample entries from the `/etc/group` file on that same system are as follows:

```
adm:x:4:wvh,juser
cdrom:x:24:hal,wvh,juser
lpadmin:x:106:wvh,juser
```

```
admin:x:112:wvh,juser
wvh:x:1000:
```

Each entry in the `/etc/group` file begins with the name of the group, a password field (in these examples indicating that any group passwords are stored in the file `/etc/shadow`), the numeric identifier associated with that group, and the list of users who belong to that group.

As you can see from these examples, it is common for users to belong to multiple groups, each of which provides access to a specific resource such as files, directories, or administrative capabilities. You can list the groups that you belong to by executing the `groups` command at any Linux command-line prompt.

When you create user accounts on an Ubuntu system (including during the installation process), a default group is created for each of those users. This group is the group of which that user is a default member when he or she logs in on your system. Most applications that use group membership to indicate administrative privileges will automatically check the `/etc/group` file when they first execute. To access shared directories, you may have to change the group that you are actively a member of. You can do this using the `newgrp` command at any Linux command-line prompt.

Ubuntu systems reserve group numbers less than 1000 for administrative purposes. The different applications that use these groups, generally referred to as system groups, automatically check group memberships when they execute. There's no need to memorize the default groups used by different system applications on Ubuntu because the administrative applications that require them typically add you to the right group. See the section "Performing Privileged Operations in Ubuntu" later in this chapter, for details.

File and Directory Permissions Under Linux

This section provides an introduction to how file and directory permissions are displayed and used on Linux systems. Although discussed in more detail in Chapter 14, "Using Command-Line Tools," it makes sense to introduce this topic here to illustrate how user and group identities can provide shared access to files, directories, and other resources, and how to manipulate those settings.

The easiest command to use to view the permissions on a file or directory is the Linux `ls` (list) command, which displays information about the files and directories in a specified location.

For example, the command `ls -ld /home/wvh` provides a long listing of my home directory, which includes information about the current permissions on that directory:

```
drwxr-xr-x  145 wvh wvh 7728 2006-03-05 10:54 /home/wvh
```

I'll focus on the first field, which shows the current permissions on my directory — see Chapter 14, "Using Command-Line Tools," for detailed information about all of these fields, options to the `ls` command, and the `ls` command itself.

The permissions field of `ls` output can be broken down into four sections:

- The first character, which identifies the type of object you're looking at. The most common of these is a `-` (a dash) if the thing you're listing is a regular file, a `d` if it's a directory, a `c` if it's a device node that can be accessed as a stream of characters, or a `b` if it's a device node that can be accessed as a block device.
- Three sets of three characters, which represent the permissions that the owner, the group owner, and all others have on the file or directory.

The most common values for each position in the three permissions sections are `r`, which means that the file or directory can be read; a `w`, which means that the file or directory can be written to; and an `x`, which (for a file) means that the file can be executed or (for a directory) means that the directory can be searched for other files or directories. If any of these permissions are not set, its position is represented by a dash. In addition, the user and group execution bits can be set to an `s`, which means "set user or group ID upon execution." In other words, executing that file is done as though it were being done by the owner and/or group of the specified file. This is commonly done to execute a command as though it were being executed by another, more privileged user on your system.

Performing Privileged Operations in Ubuntu

Aside from your own user name and numeric ID, the most important user name on a Linux system is the user named root, whose user ID and group ID are both `0`, and who is often known as the *superuser*.

On most Linux systems, privileged operations are often done by using the `su` (substitute user ID) command to become the root user. However, Ubuntu does things slightly differently. Ubuntu uses the `sudo` (substitute user ID do) command to perform all privileged operations. The `sudo` command uses the text-format configuration file `/etc/sudoers` to determine which users can perform privileged operations as the superuser. On Ubuntu systems, any member of the admin group can perform privileged operations as the root user. You can use the `sudo` command on any Linux system, but you cannot use the `su` command on a Ubuntu system.

The difference between using the `su` and `sudo` commands is subtle but significant:

- When using the `su` command to perform a privileged operation, you execute the `su` command and supply the root user's password in response to the password prompt, which then starts a subshell with root privileges. You then execute whatever privileged commands you want within the context of that shell. They are all therefore executed as the superuser. When you are done, you can either exit from that shell or suspend it for subsequent reuse.
- When using the `sudo` command, you execute the `sudo` command, followed by the name and arguments to the command that you want to execute.

NOTE **The `su` and `sudo` commands are not restricted to executing commands as the root user. You can become any user using the `sudo` command by specifying that user's name after the `su` command and providing that user's password in response to the password prompt. Similarly, you can use the `sudo` command's `-u` option to specify the name of the user that you want to execute a command as. As always, the `su` command requires that you know a specific user's password, while the `sudo` command requires only that you provide your password.**

Let's look at a few examples to compare using these commands. The following is an example of using the `su` command to display the partitions available on a disk:

```
$ su
Password: <enter-root-password>
#  fdisk -l /dev/sda
[output]
# exit
$
```

To do the same thing using the `sudo` command, you would do the following:

```
$ sudo fdisk -l /dev/sda
```

```
Password: <enter-your-password>
[output]
$
```

On Linux systems other than Ubuntu, you could also use the `su` command's `-c` option (execute a single command) to do something similar, as in the following example:

```
$ su -c fdisk -l /dev/sda
Password: <enter-root-password>
[output]
$
```

The key difference between the `su` and `sudo` command is whether the command prompts you for the root password or your personal password. However, Ubuntu's focus on using the `sudo` command has a few basic advantages:

- The system is harder to attack because an attacker first has to discover the identity of a privileged user on the system before they can attempt to break in.
- Users have to remember only their personal password to perform privileged operations if they are permitted to do so. There is no need to separately set and secure the system's root password.

TIP **If you want to execute several privileged commands in a row on a Ubuntu system, you have two alternatives. First, you can execute the `sudo -s` command to execute a shell as the root user. Second, you can do the same thing by running the `sudo /bin/bash` command. In both of these cases, you have to remember to explicitly exit from the root shell once you're done running the commands that you want to execute.**

Summary

This chapter covered some of the basic concepts of Linux and Ubuntu. The chapter began by introducing files and directories and explaining the directories that are commonly used on Ubuntu systems. It then discussed how Ubuntu uses disks and partitions them into separate filesystems. I covered the different types of filesystems that you can use on partitions or on Ubuntu systems in general, and how to make those filesystems available for use on your system. The last few sections introduced basic Linux administrative concepts such as managing users, working with groups and file permissions, and executing privileged commands on the Ubuntu system. This chapter provided the basic knowledge that you'll need to understand these concepts when you encounter them elsewhere in this book.

Chapter 5

Using the GNOME Desktop

Although many curmudgeons and long-time UNIX users reject any sort of graphical interface, let's face it — most people today want (and expect) one. The graphical environment used on Ubuntu systems, the GNOME desktop, provides a stable and usable environment for running your graphical applications and interacting with your system graphically. Most of the Linux utilities used for system administration and configuration provide graphical interfaces to simplify formerly complex tasks, and are easily accessed from one of the primary menus provided by the GNOME desktop.

This chapter begins by providing some background information on the graphical environment used on all Linux systems, such as explaining exactly what the word "desktop" means, and what graphical alternatives exist on Linux systems. This information provides a good foundation for understanding how graphical interfaces work on Linux systems, regardless of the specific graphical environment that you choose to use. The majority of the chapter focuses on the organization and use of the GNOME desktop provided on standard Ubuntu desktop systems. GNOME is a powerful graphical interface with all of the features that you'd expect in a modern graphical user interface (GUI) — once you know where to find them. The final section explains how you can further customize GNOME by replacing the software that it uses internally to manage window placement, movement, and so on. If you are using Kubuntu, you should still find the introductory information in this chapter to be useful, although you may want to focus on Chapter 7, "Using the KDE 3 Desktop," or Chapter 8, "Using the KDE 4 Desktop," depending upon the version of Kubuntu that you are using. These chapters provide information about using KDE, the graphical desktop environment that is used by the Kubuntu edition of Ubuntu.

IN THIS CHAPTER

- Using the mouse
- GNOME basics
- Using the desktop
- Personalizing the desktop
- Using the Nautilus file manager
- Desktop search on Ubuntu
- Using a window manager

What's a Desktop? Graphical Environments for Linux

Almost all of the high-resolution support for interacting with any Linux system is handled by a graphics package called the X Window system, known to its friends as X11 or simply X. The X Window system is one of the most attractive aspects

of Linux and almost any operating system running on modern bitmapped graphics workstations. X is a network-aware graphics windowing system that provides a similar set of capabilities for creating and working with applications as do the graphical environments used on Microsoft Windows and Mac OS X systems. The X Window system was designed to provide a common windowing environment on multi-processing, networked computer systems and is the industry standard windowing system for computers running different versions of Linux, any other UNIX-like operating system, and even computers that are still running Compaq's VMS operating system. Versions of the X Window system are also available for all releases of Windows greater than 3.1, and for any version of Mac OS X — these run X Window system applications in the context of the native windowing systems for those platforms.

The X Window system was originally developed at the Massachusetts Institute of Technology (MIT), was under the custodianship of the Open Group (`www.opengroup.org`) for a while, and is now stewarded by the X.org Foundation (`www.x.org`). A previous implementation of the X Window system for Linux systems, known as XFree86, is no longer used on Linux systems — the official source of the X implementation for Linux systems is now X.org.

The X Window system enables you to run multiple applications on a bitmapped display screen, each of which displays its own window(s), graphics, dialog boxes, and so on in one or more separately controllable windows on the screen. An X Window system window manager is an application that also runs on a bitmapped display screen, but its job is to manage the individual windows that are created and used by other applications. The window manager is the application that enables you to move windows around on the screen and raise and lower windows. The window manager displays the window borders that enable you to move applications; displays the menus that pop up when you click on a part of the screen that is not occupied by a window; displays the window controls that enable you to minimize, maximize, and close separate windows; and so on. Popular window managers include After Step, Black Box, Compiz, Fluxbox, FVWM, ICEwm, kwm, Metacity, twm, vtwm, and Window Maker. All of these provide the same basic functionality, although they differ in terms of how and where they manage minimized applications, how the window manager is configured, how applications are bound to menus and/or graphical areas on your screen, how they support alternate fonts and display styles, and so on.



One great feature of the X Window system is that, because it was designed to work across multiple hardware platforms, you can start X Window system client applications on one system but have them display on the bitmapped screen of another system (assuming that this capability is not blocked by a firewall or system configuration setting). This is done by setting an environment variable called `DISPLAY` and then starting the application that you want to display remotely. The `DISPLAY` environment variable has the following form:

```
host:display.screen
```

Here, *`host`* specifies the name or IP address of the system that you want to display the window on, and is followed by a colon. After the colon, the first number identifies the graphics *display* on which the window should be displayed, followed by a period, and the number followed by a colon, followed by the number of the physical screen attached to that card. In most cases, both `display` and `screen` are `0`.

See the section of Chapter 14 entitled "Using Environment Variables" for more information about setting environment variables. In many cases, you can also start an X Window system application on a remote screen by starting the application from the command line and specifying the `-display host:display.screen` argument.

After window managers, the next step in the evolution of the X Window system desktop experience is the desktop manager. Desktop managers, colloquially referred to simply as desktops, always run a window manager under the covers, but also provide additional capabilities such as the following:

- Drag and drop
- A file manager of some sort to simplify graphically browsing files and directories
- Easy, flexible ways to tie applications to buttons, icons, and menu items
- A centralized mechanism for configuring the appearance of your screen and the windows that it displays, generally referred to as applying a *theme*
- Support for running lightweight applications (known as *applets*) within the context of the window manager to support certain types of tasks

The most common desktops for Linux systems are GNOME (GNU Network Object Model Environment, `www.gnome.org`), KDE (K Development Environment, `www.kde.org`), ROX (`rox.sourceforge.net`), and XFCE (`www.xfce.org`). Both GNOME and KDE are rich, robust desktop environments that provide their own window managers (on Ubuntu systems, GNOME uses Compiz, while KDE uses kwm) and primarily differ in terms of default mouse and key bindings, menu organization, and the underlying graphical toolkit. (GNOME uses a widget toolkit known as GTK, while KDE uses one called QT.) For some background information on the relationship between GNOME and KDE, see Chapters 7 and 8. Both ROX and XFCE are lightweight desktops, which means that they attempt to strike a balance between the heavy resource requirements and corresponding power of rich desktop environments such as GNOME and KDE, and the much lighter requirements but more limited capabilities of traditional window managers. You can find a great introduction to most of the X Window system window managers and desktop environments available for Linux at `http://xwinman.org`. Versions of Ubuntu prior to the 7.10 (Gutsy Gibbon) release used the Metacity window manager with GNOME, but versions 7.10 and later use the Compiz window manager. For detailed information about the Compiz window manager, see Chapter 6, "Using the Compiz Window Manager."

The rest of this chapter primarily focuses on GNOME, which is the desktop environment that is used on Ubuntu systems, and which starts automatically the first time you log in to a Ubuntu system. Although you can change the window manager used by GNOME, switch to KDE, or even switch your Ubuntu experience to using a window manager rather than using its default GNOME configuration, those are all advanced tasks. First things first — let's explore how things work on a standard Ubuntu desktop system just in case your goal is getting your work done rather than tweaking your Ubuntu system. There's plenty of time in the future to trick out your graphical environment, which I'll discuss in the section entitled "Switching to a Window Manager" later in this chapter.

Favorite desktop environments are like opinions — everybody has one. Even those of us who have used many different window managers and desktops over the years, and are therefore reasonably agnostic, have our own favorites. If you prefer KDE or Kubuntu (which uses the KDE desktop environment) or KDE was recommended to you, see Chapters 7 and 8, entitled "Using the KDE 3 Desktop" and "Using the KDE 4 Desktop" respectively, (depending on the KDE version that you want to use) for details about using that desktop.

Using the Mouse

Interacting with a graphical user interface of any kind traditionally requires the use of a mouse and keyboard to control the cursor, to select and move items, access menus, and enter data within applications. As you'll see throughout this book, the actions associated with the different mouse buttons do different things depending on the portion of the screen or application in which you use them. This is known as being *context sensitive*, and is one of the keys to performing many different types of actions using only three little buttons.

By default, the mouse in GNOME is configured for a right-handed user, with the different mouse buttons assigned to the following actions:

- **Left button:** Single-clicking selects an icon, menu, menu item, or object within an application. Double-clicking on a desktop icon opens the item associated with that icon. Holding down the left button and dragging the cursor in an application or on the desktop selects multiple icons or objects within that application. Holding down the left button and dragging a selected object moves that object. These operations can be (and often are) combined. For example, in a word-processing program, you frequently left-click once to position the cursor, hold down the left button and drag the cursor to select a block of text, and then left-click, hold down the mouse button, and drag the selected text to a new location in your text.
- **Middle button:** Single-clicking pastes text or objects that have been cut or copied, or moves a selected window behind other windows that are currently displayed on the screen. Middle-clicking on a desktop icon and dragging it to a new location displays the Move menu, which enables you to move or copy icons to new locations without actually relocating the icon until you confirm your intentions.
- **Right button:** Displays any context-sensitive menu that is available at the current cursor position.

You can configure how reactive the mouse is, what your mouse cursor looks like, and what the various mouse buttons do using the Mouse Preferences application, which is explained later in this chapter in the section "Customizing Mouse Behavior." For example, if you are not right-handed and haven't been beaten into submission so that you use the mouse with your right hand, you can configure the mouse for left-handed use (which reverses the meaning of the various buttons) using this application.

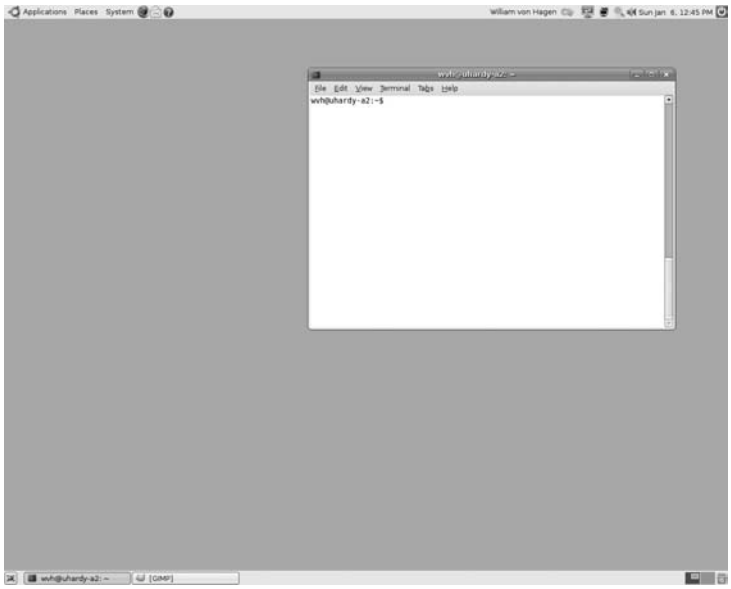
NOTE **If you have trouble seeing the screen or using a mouse or keyboard because of a physical condition, don't worry — the GNOME and Ubuntu folks are sensitive to your needs. See the section "Assistive Technologies for Using GNOME" later in this chapter for more information. The KDE folks are similarly attentive — additional information on assistive technologies for KDE is provided in Chapters 7 and 8.**

GNOME Desktop Overview

Figure 5-1 shows the default GNOME desktop on an Ubuntu Linux system the first time you log in. This figure also displays a single application window as an example — you won't see that unless you explicitly select the Applications Accessories Terminal menu command, but I wanted to be able to explain the window controls that are available in any GNOME application window. The Ubuntu folks have gone through a lot of effort to create an attractive, eye-pleasing background, set of fonts, window decorations, and control buttons for any applications that you start on your Ubuntu system. These window decorations and controls are known as a *theme*, which you can easily customize yourself, as explained later in this chapter.

FIGURE 5-1

The default GNOME desktop in Ubuntu Linux



NOTE **Figure 5-1 shows the default desktop — in other full-screen screenshots, I've changed the background to a solid color to make things easier to see. Changing the default background is discussed in the section "Changing Desktop Backgrounds," later in this chapter.**

Figure 5-1 identifies the various portions of the default GNOME desktop. The following describes them in more detail:

- **Top Panel:** A panel is a special portion of your desktop, controlled by the gnome-panel application, which provides access to various ways of interacting with the desktop and launching different types of applications. GNOME can display panels along any edge of the screen, but the default GNOME configuration displays panels at the top and bottom of the screen. By default, the top panel contains the following items:
 - **Applications, Places, and System Menus:** A menu consists of a list of shortcuts to specific commands, tasks, or applications. These three menus provide easy, graphical access to applications, portions of your computer system or network locations, and system-related tasks, respectively. Each of these menus is described in more detail in the next section, "Menus in GNOME." To display any of these panel menus, you simply left-click on the name of the menu, and the associated menu drops down and displays on your screen. To close the menu, simply left-click elsewhere on the screen.
 - **Application Launchers:** These icons start specific applications for you when you left-click on them.
 - **Notification Area:** A portion of the panel where applications and applets display icons in response to specific events, such as receiving mail, low-power conditions, and so on. One

especially important icon that can be displayed in this area is the Updates icon, which alerts you that updates to your system or the packages that you've installed are available.

- **Applets:** These icons launch lightweight applications that run in the context of a panel and provide various capabilities. Adding and configuring applets is discussed later in this chapter in the section entitled "Customizing Panels."
- **Log Out Button:** This button quickly terminates your current login session and re-displays the GNOME login manager.

- **Workspace:** This is the portion of the screen in which applications execute and display associated windows and dialogs.
- **Bottom Panel:** This panel is displayed at the bottom of the screen and contains the following items by default:
 - **Show Desktop Button:** This button minimizes all windows and dialogs that are currently displayed on the screen, revealing the current desktop.
 - **Application Controls:** Each application that is active on your current workspace displays an associated control region in this portion of the bottom panel. Right-clicking on this control alternately minimizes and maximizes the application, while left-clicking on this control displays a context-sensitive menu for moving and controlling that application's window(s).
 - **Workspace Switcher:** This is a special applet that enables you to manage multiple workspaces and provides a miniature display of each active workspace. Workspaces are essentially separate virtual screens that are provided by the GNOME desktop. You can run applications on different workspaces, move applications between workspaces, and so on. Multiple workspaces provide a convenient way of running different types of applications on different virtual screens without them being visible until you actually switch to the workspace where they are displayed. This is useful, for example, when you're at work and want to play a game — you can start the game on a different workspace from the one where you're actually working, and switch to it when no one is looking.
 - **Trash Applet:** This is a special applet that provides access to the Trash Can, which is a special portion of the GNOME desktop to which you can drag files to subsequently delete them.

As you'll see throughout the rest of this chapter, GNOME is extremely configurable — the chances are that no two users' GNOME desktops will look the same, which is completely fine. How you configure and use your computer system is up to you, and should reflect your interests and the types of things that you want to do. The job of desktops such as GNOME is to make it as easy and as comfortable as possible to accomplish whatever task you are trying to do.

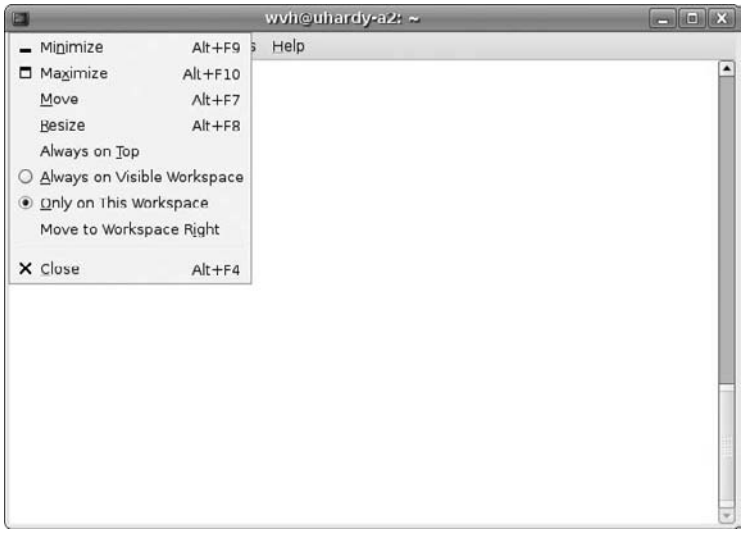
GNOME Application Windows

The GNOME application window shown in Figure 5-1 provides a number of window controls that enable you to move the window around on the screen, control its size, and so on. The window manager that you are using determines the location of these controls, the number of controls displayed, and the capabilities provided by these controls. These window controls are shown in the top part of the frame around the application window (typically known as the title bar) in most window managers.

The control at the far left of the title bar displays a menu of window control commands. The contents of this menu depend on the window manager that you are using. The window control menu displayed by the default window manager used by GNOME on Ubuntu systems, Compiz, is shown in Figure 5-2.

FIGURE 5-2

The Compiz window manager's window control menu



The three window controls at the right of the title bar in Figure 5-2 do the following, from left to right:

- Minimize the current window
- Switch between a full-screen or a user-specified window size
- Close the current window, terminating any processes that it is running

All of these window controls are fairly standard, and are found in the window decoration and title bar layouts used by most modern window managers.

Menus in GNOME

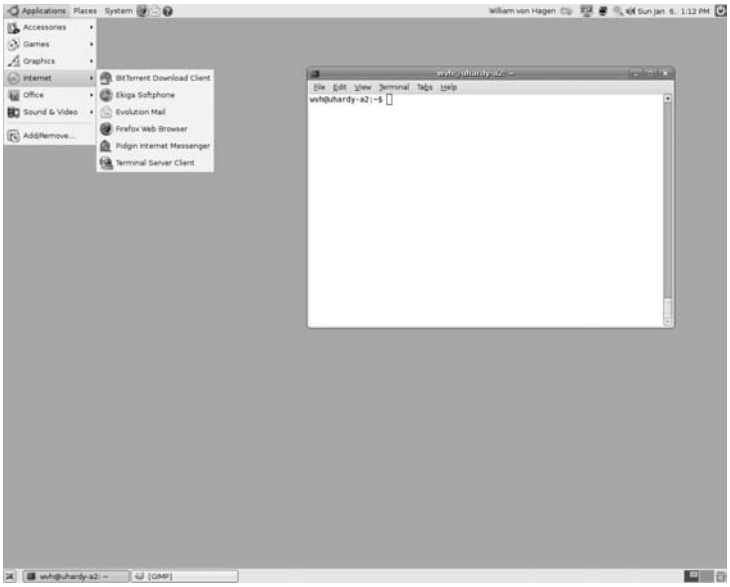
As discussed earlier in this chapter, GNOME provides two different types of menus: menus located in a panel, which provide easy access to different applications, locations, and capabilities on your system, and context-sensitive menus that provide access to different capabilities depending on the context (application, desktop, and so on) in which you've activated them. This section provides a more detailed discussion of these different types of menus — how to display them and how to customize them, and the types of actions that they provide access to.

Panel Menus

Menus are displayed in the panels when you left-click on the menu name in the panel. Menus either contain entries for specific tasks or locations on your system, or contain other menus. Menus located on other menus are known as *submenus*. As an example, Figure 5-3 shows the default contents of the Application menu's Internet submenu, which I discuss later in this section.

FIGURE 5-3

Menus and submenus in GNOME



The next few sections describe the standard panel menus and their contents in more detail.

The Application Menu

The Application menu makes it easy to launch different types of GNOME applications, grouping them into submenus that display when you move the cursor over each menu entry. The Application menu provides the following groups:

- **Accessories:** Contains a general category for small, commonly used applications such as a calculator, dictionary, text editor, and command-line terminal.
- **Games:** Contains exactly what you'd expect. Ubuntu is designed for people, so it naturally includes several games as part of its default installation. For more information about games on Ubuntu Linux, see Chapter 20, "Would You Like to Play a Game?"
- **Graphics:** Provides access to several applications for creating graphics, and capturing and manipulating images off the screen or from a scanner, and a small, lightweight image viewer for pre-

viewing existing graphics files. See Chapter 18, "Working with Graphics," for more information about graphics software on Ubuntu Linux systems.

- **Internet:** Contains entries for accessing information sources found on the Internet, such as a Voice-Over-IP telephone system, the Evolution e-mail client (discussed in Chapter 9, "Managing Email and Personal Information with Evolution), the Firefox Web browser (discussed in Chapter 11, "Surfing the Web with Firefox"), instant messaging and IRC clients, and so on.
- **Office:** Lists the standard business applications available on your system. By default, this menu includes an entry only for the Evolution e-mail client, but you can easily add entries for your favorite word processors (see Chapter 16, "Creating and Publishing Documents"), spreadsheet and presentation software (see Chapter 17, "Other Office Software: Spreadsheets and Presentations"), and any other Linux software that you use as part of your daily activities.
- **Sound and Video:** Provides easy access to audio and video applications, including applications for recording sounds, converting your (officially purchased) audio CDs to online format, and playing online audio and video, and sophisticated applications for creating different types of CDs, VCDs, and DVDs. See Chapter 19, "Working with Multimedia," for more information about working with different types of multimedia files and mediums on Ubuntu Linux systems.

In addition, the Applications menu provides an Add/Remove Applications entry that makes it easy to add applications to the menu. As discussed later in this section, the Add/Remove Applications application is seamlessly integrated with Ubuntu's package update mechanism — if you add an application that is not already installed on your system but is available from the Ubuntu repositories, the application will be downloaded and installed as part of the menu update. Pretty cool!

The Places Menu

The Places menu provides quick access to files, directories, and storage-related activities. Selecting a location or storage device on the Places menu opens the Nautilus file manager, displaying the contents of the selected location. The Places menu also provides convenient access to storage-related actions, such as writing to a CD or DVD, connecting to a file server, and searching for or revisiting specific files and directories. For more information about the Nautilus file manager, see the section later in this chapter entitled "Introducing the Nautilus File Manager."

NOTE **The contents of the Places menu on your system will change as you add and remove external storage devices to your system, visit remote file servers, and so on.**

Aside from shortcuts to specific locations on your system, the Places menu provides a Recent Documents submenu, which contains a list of the ten files or desktop objects that you've most recently opened. This menu makes it easy for you to continue working on the same documents each time you log in on your Ubuntu Linux system.

The System Menu

The System menu provides access to commands for interacting with GNOME and your Ubuntu system, such as setting personal preferences for how the GNOME desktop looks and behaves, performing system administration tasks such as configuring hardware, updating your system, and configuring network services, getting help, and logging out of or shutting down your Ubuntu system.

The System menu contains the following items:

- **Preferences:** Provides access to a large number of applications that enable you to configure how you interact with GNOME, your Ubuntu system, and any networks that you have access to (such

as the Internet or a home network). If you don't like the way that something works on your Ubuntu system, check the items on the Preferences menu for something that seems appropriate, start any applications that appear to be relevant, and explore a bit. The section "Customizing Your Desktop," later in this chapter, explains how to perform many of the most common GNOME and Ubuntu configuration tasks using commands from this menu.

- **Administration:** Provides access to applications that enable you to configure the software installed on your Ubuntu system, the devices attached to your Ubuntu system, the services that your system runs, and core system information such as the time, the date, the list of users who are authorized to access your system, and so on.
- **Help and Support:** Provides different types of help for using your system. This menu provides access to the online documentation installed on your Ubuntu system, a shortcut to the Ubuntu documentation team's Web site on the Internet, and even a shortcut to the Ubuntu Marketplace (discussed in Chapter 1), where you can identify firms that may be able to help you with any tough problems or configuration issues that I may have missed in this book.
- **About GNOME:** Displays information about the specific version of GNOME that you are running, as well as general information about the many wizards who have contributed to the GNOME project in one way or another. You will probably need to provide GNOME version information if you post problem reports or questions, or if you contact any of the Ubuntu vendors in the Ubuntu Marketplace.
- **About Ubuntu:** Displays online documentation for the Ubuntu Project, which includes information about the specific version of Ubuntu Linux that you are using. As with GNOME version information, you will probably need to provide Ubuntu version information if you post problem reports or questions, or if you contact any of the Ubuntu vendors in the Ubuntu Marketplace.
- **Quit:** Provides a convenient shortcut to an application that enables you to log out of the system and can optionally cleanly shut it down and turn it off for you after performing several housekeeping tasks. See Chapter 22, "Understanding the Ubuntu Startup and Shutdown Processes," for more information about how your system boots and what it does when it shuts down.

Context-Sensitive Menus

The menus discussed in the previous section always display when you left-click on a specific menu heading or submenu, regardless of what else is happening on your Ubuntu Linux system. Context-sensitive menus display when you right-click on something on your Ubuntu Linux screen, but their contents depend on what you've right-clicked on. Because their contents differ based on the context in which you've right-clicked, these right-click menus are referred to as context-sensitive menus. Figure 5-4 shows the context-sensitive menu that you see when you right-click on any portion of the background for your Ubuntu desktop.

To show how application-specific these different context-sensitive menus can be, Figure 5-5 shows the context-sensitive menu that displays when you right-click within the main window for the Firefox Web browser.

FIGURE 5-4

The GNOME desktop's context-sensitive menu

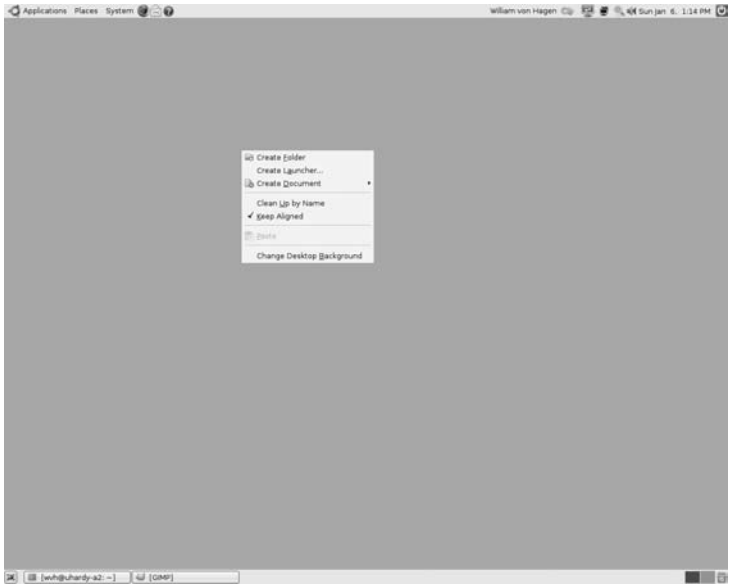
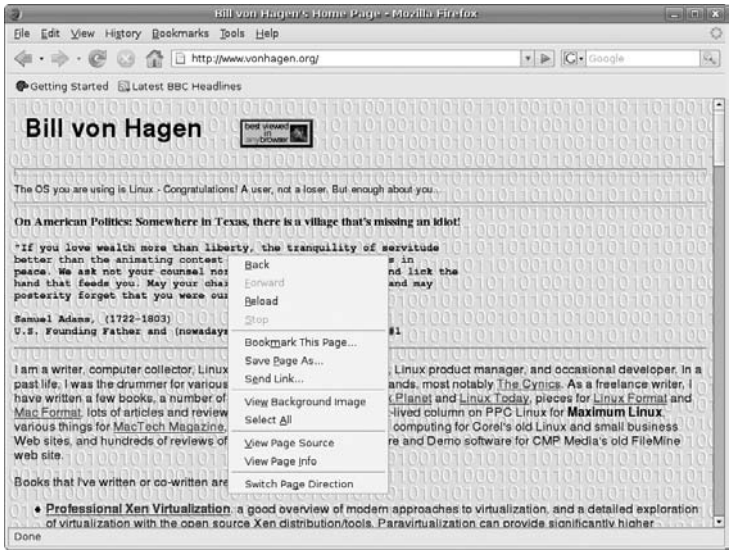


FIGURE 5-5

A context-sensitive menu in the Firefox Web browser



The contents of context-sensitive menus are application-specific, which means that you can't directly customize them. However, these menus can be customized by installing additional applications or applets, or (in the case of extensible applications such as Firefox or Nautilus) by installing extensions or plug-ins that add new capabilities and therefore provide associated menu items.

Customizing GNOME Menus

Ubuntu Linux provides you with two ways to customize the contents of the basic GNOME menus. The Add/Remove item at the bottom of the Applications menu makes it easy for you to add new applications to any of the pre-existing sections defined on the Applications menu by installing those applications and creating their menu entries as part of the installation process. Using the Add/Remove menu item is discussed in the section of Chapter 23 entitled "Installing Software with the Add/Remove Menu Item." This menu item is extremely convenient if you want to add some popular application to the Applications menu, but isn't much use if you want to add your own menus or create custom menu entries for specific applications. To do this, you'll need to use the System Preferences submenu's Main Menu application. Selecting this menu item starts the Alacarte Menu Editor, which makes it easy for you to add almost anything (including custom submenus) to any existing GNOME menu.

For example, suppose you wanted to create your own submenu called Favorites and begin adding entries for your favorite applications to it. To do this (adding the Emacs text editor to this menu as an example), do the following:

1. Select System Preferences Main Menu to start the menu editor. This menu editor dialog consists of two primary panes: a menu catalog pane at the left and a menu contents pane at the top right that shows the current contents of the currently selected menu. The Applications menu is expanded when this dialog opens, as shown in Figure 5-7.

TIP **You'll notice that the menu catalog pane lists many menus that are not currently displayed on your system. This is because they are not marked to be shown (via the Show column) in the menu contents pane on the right, as shown in Figure 5-6.**

FIGURE 5-6

The Alacarte Menu Editor on your Ubuntu system



2. Click the New Menu button. A dialog like the one shown in Figure 5-7 displays (with no entries in the text fields).

FIGURE 5-7

Adding a new menu to your Ubuntu system



3. Enter the name of the new menu that you want to create (I'm using Favorites in this example), add an optional comment that will be used as the tool tip for this menu, and click OK. The new menu will be added to the Applications menu, and the Applications menu entry in the menu catalog pane will expand to show the current contents of that menu (which should be identical to the list in the menu dialog at the top right). Select the name of your new menu in the menu catalog pane. The contents of the new menu (i.e., nothing) displays in the menu contents pane at the top right, as shown in Figure 5-8.
4. Select the File New Entry button. A dialog like the one shown in Figure 5-9 displays.

FIGURE 5-8

The new menu added to your Ubuntu system

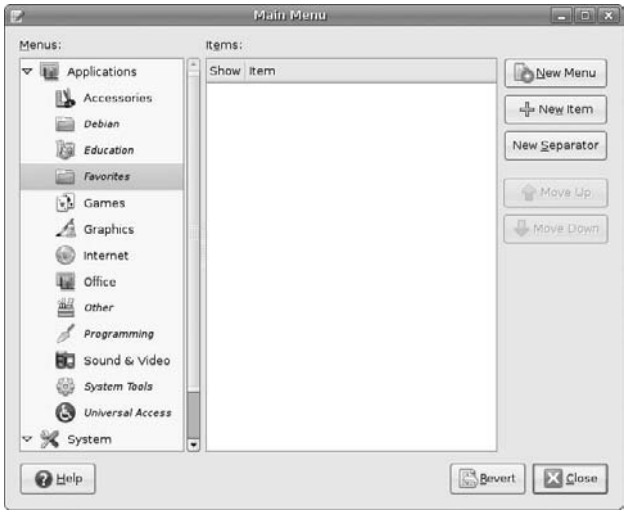


FIGURE 5-9

Adding a new menu entry to your Ubuntu system

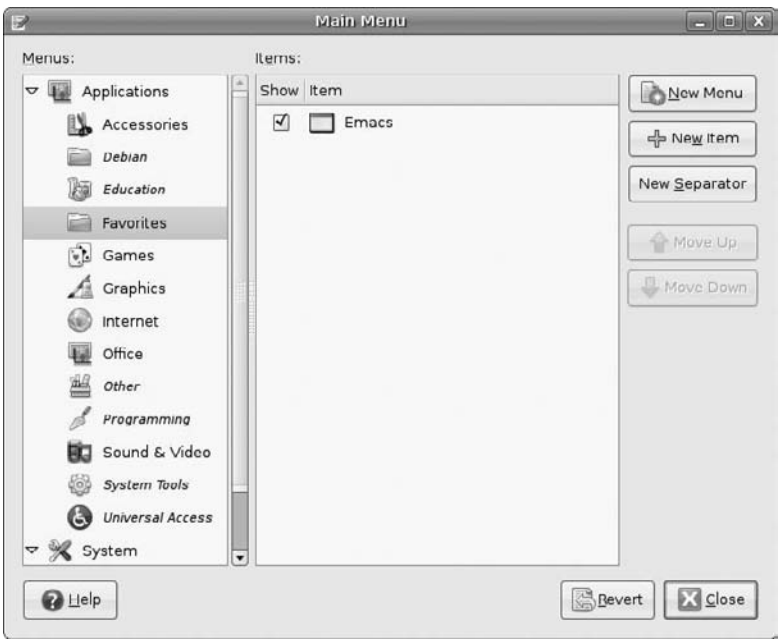


5. Enter the information about your new entry (in this case the Emacs text editor) and click OK. The dialog closes and the new entry on your Favorites menu displays, as shown in Figure 5-10.

TIP **If you don't know exactly where the application that you want to add is located, you can click Browse to display a browse dialog that lets you navigate to and select the application for which you want to create a new menu entry.**

FIGURE 5-10

Your new menu entry in the menu editor



6. Select the Applications menu in the menu catalog pane at the left, and make sure that the checkbox in the Show column for your new menu is selected in the menu listing pane at the top right, as shown in Figure 5-11. This entry is required so that the menu is listed in the Applications menu once you exit from the menu editor.

FIGURE 5-11

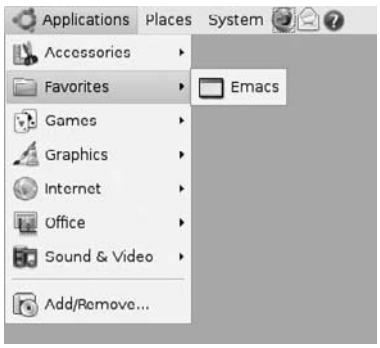
Making sure that the new menu is visible



7. Click Close to exit from the Alacarte Menu Editor and add your changes to the Applications menu. The Applications menu is updated, and your new menu and entry are now available, as shown in Figure 5-12.

FIGURE 5-12

Your new entries on the GNOME Applications menu



If you decide not to save your changes, you can click the Revert button to undo your current changes, and then click Close.

The flexibility provided in the dialog shown in Figure 5-9 makes it easy for you to create menu entries for applications that you have manually compiled and installed yourself, such as in corporate or academic environments, and also to add your own favorite command-line parameters to new or existing menu entries.

TIP **You can also add menus in GNOME using the Nautilus file manager, but I find that using applications designed for editing menus is easier to understand. For information about using Nautilus to create menus and menu items, see the online GNOME Users Guide, available through the System menu's Help menu.**

Customizing Your Desktop

As explained earlier in this chapter, GNOME makes it easy to customize its menus and the sections and applications that they contain. Luckily, every part of the GNOME desktop is similarly configurable. This section highlights how to do some of the most common customizations that most people do to make their systems easier to use and generally personalize them. Many other customizations are available, but discussing all of them could make this book impossible to lift. Let's focus on the highlights!

Customizing Mouse Behavior

One of the most common customizations that people want to do is to change how responsive the mouse is, in one way or another. Available customizations include:

- Whether the mouse is configured for left-handed or right-handed users. The latter is the default, but may not suit you if you use your left hand to control the mouse.
- The speed at which the cursor moves and the amount of screen real estate the mouse covers in relation to your moving the mouse
- The period of time that can elapse between two mouse clicks to consider them as a double-click (which opens the target) rather than simply two single clicks (which simply selects something). If you have difficulty double-clicking, you can increase the double-click time period.

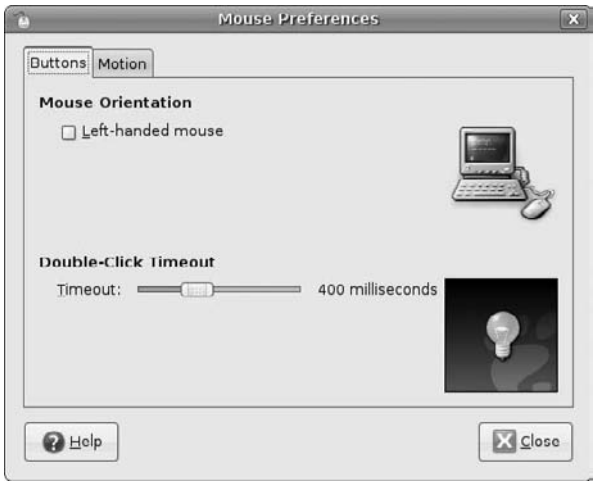
All of these customizations are controlled by the System Preferences menu's Mouse application, which is shown in Figure 5-13.

This application provides two tabs that enable you to configure different aspects of the mouse: the behavior of the buttons and the speed and sensitivity of mouse motion. The Buttons tab is shown in Figure 5-13. To switch between different tabs, simply select the appropriate tab in the Mouse Preferences application.

The configuration items in these tabs are pretty straightforward, so I won't explain each of them. However, the tab that does require some explanation is controlling the motion and sensitivity of the mouse. These options are located on the Mouse Preferences application's Motion tab, which is shown in Figure 5-14.

FIGURE 5-13

The Mouse Preferences dialog



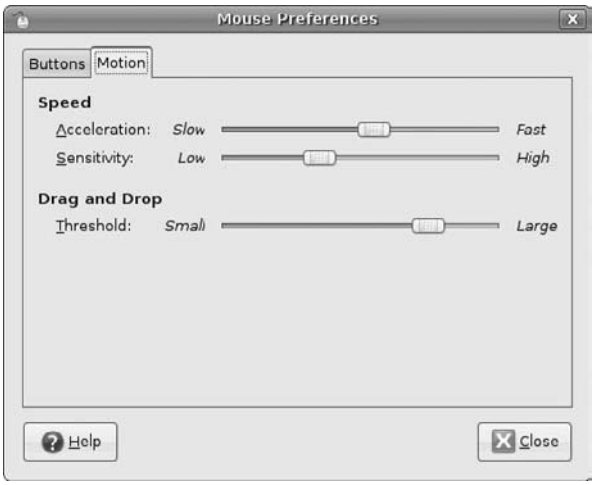
The entries on this tab include the following:

- **Acceleration:** This option controls how far the cursor moves on the screen in response to your moving the mouse, and is controlled by a slider. Being able to change this ratio guarantees that you don't need a mouse pad the size of your computer monitor — acceleration enables the cursor to cover larger distances as you move the mouse continuously. Slower values mean that cursor movement is closer to the actual distance that you're moving the mouse; faster values mean that the cursor moves proportionately farther distances as you continue to move the mouse — in other words, the cursor starts out moving slowly and accelerates as you continue to move the mouse. To change your settings, click and hold the button on the slider using the left mouse button and drag it to the new value that you want to use.
- **Sensitivity:** This option determines how sensitive the cursor is to any movement of the mouse. If you have set Acceleration to be very fast, you may find that the cursor jumps when you simply touch the mouse. In this case, you could decrease the sensitivity of the mouse, while still retaining fast acceleration.
- **Drag and Drop Threshold:** This option controls the distance that the mouse must cover when you have selected a desktop object in order for the motion to be considered a drag-and-drop action. For example, when double-clicking, you may find that you move the mouse slightly between the two clicks. If the drag-and-drop threshold is set very low, GNOME will interpret this as a drag-and-drop or move operation rather than a double-click operation. In this case, you would want to increase the threshold value.

Changes to your mouse preferences take effect immediately, so it's easy for you to experiment with the set of values that gives you the control and behavior that you want on your system.

FIGURE 5-14

Configuring the mouse motion



Configuring Screen Resolution

When you installed Ubuntu, one of the questions you were asked was to select the display resolutions that you want to be able to use. It's nice of the Ubuntu installer to ask this question, but display resolution is one of those things that's hard to visualize in advance — it really makes sense only when you can see what different resolutions actually look like. Consequently, even if you are using a 30-inch monitor and the sexiest video card in the universe, you may find that your Ubuntu Linux system insists on displaying everything at a resolution such as 1024 768, where each of the icons on your screen is approximately the size of a grapefruit. This is probably not what you had in mind, and is certainly not the best use of your hardware or screen real estate. This section explains how to correct that problem by selecting a higher screen resolution. Internally, this means that your screen displays more pixels, but the visual effect is that things on the screen are both smaller and sharper. In the computer industry, this is known as "the right thing™."

Screen resolution in GNOME is controlled by the System Preferences menu's Screen Resolution application, which is shown in Figure 5-15.

Once this application displays, you can see the list of available screen resolutions by clicking and holding the left mouse button on the current resolution value. A pop-up list displays, as shown in Figure 5-16.

After selecting another screen resolution, click Apply to try out the new resolution. GNOME will change your screen to the new resolution and display the dialog shown in Figure 5-17, asking if you want to keep the new screen resolution or permanently switch to the new resolution.

NOTE **Because it's possible to select a resolution that may not actually work correctly, your screen will automatically revert to its former resolution in 20 seconds. No need to panic if your screen goes haywire for a little while!**

FIGURE 5-15

The Screen Resolution application

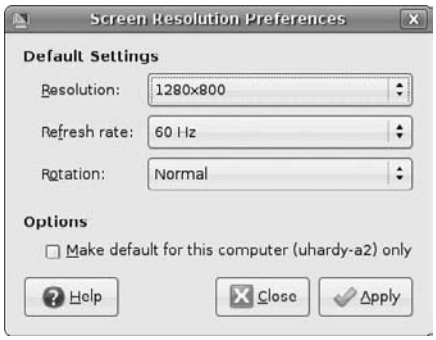


FIGURE 5-16

Selecting other screen resolutions

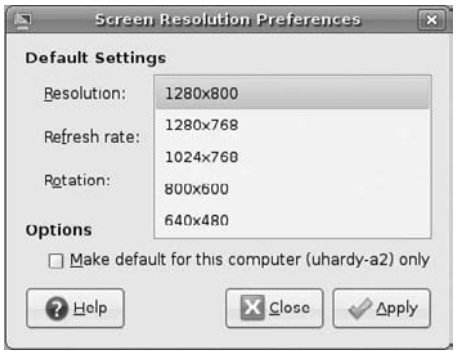


FIGURE 5-17

Confirming screen resolution changes



If you're happy with the new resolution, click Keep Resolution. The dialog will close, your screen will remain at the new resolution, and it will be the default screen resolution used whenever you log in to your Ubuntu system.

If the new resolution looks odd or is simply too small for your taste, click Use Previous Resolution. You can then experiment with other resolutions, clicking Apply to test out any newly selected resolution until you find one that you're happy with.

Customizing Panels

As discussed earlier in the chapter, panels are regions on your desktop that are displayed by the gnome-panel application and provide access to various ways of interacting with the desktop and launching different types of applications. The gnome-panel application can display panels along any edge of the screen, but the default Ubuntu configuration displays panels at the top and bottom of the screen.

NOTE **A single instance of the `gnome-panel` application controls all of the panels that are displayed on your system. The fact that multiple panels are displayed on your screen does not mean that multiple copies of the `gnome-panel` application are running.**

By default, the GNOME panels displayed at the top and bottom of your screen contain the menus and other items discussed in the "GNOME Desktop Overview" section earlier in this chapter. Not only can you configure how and where the panels are displayed on your screen and how many panels are displayed, but you can also extensively customize the contents of each. Ubuntu's default installation of GNOME comes with many other applets, buttons, and controls that you can install in the panels on your system.

Right-clicking on any existing panel displays that panel's context-sensitive menu, as shown in Figure 5-18.

FIGURE 5-18

The gnome-panel application's context-sensitive menu



Along with standard entries to display online help and get information about the version of the application, the context-sensitive menu for the gnome-panel application displays the following configuration-related entries:

- **Add to Panel:** Enables you to add a new applet, button, or other supported object to any panel that is currently displayed on your screen. See the next section, "Customizing Panel Contents," for details on adding objects to an existing panel.
- **Properties:** Displays a panel properties dialog that enables you to set certain properties for the currently selected panel. See the next section, "Customizing Panel Properties," for details on the types of things that you can customize for an existing panel.

- **Delete this Panel:** Deletes the currently selected panel from the screen and removes all information about its current configuration from the system.
- **New Panel:** Creates a new panel on the next available edge of the screen. You can create an essentially infinite number of panels — each new panel will simply start on the next edge of your screen, adjacent to any existing panel that is already associated with that edge. I find two panels to be the right number, but you may feel differently.

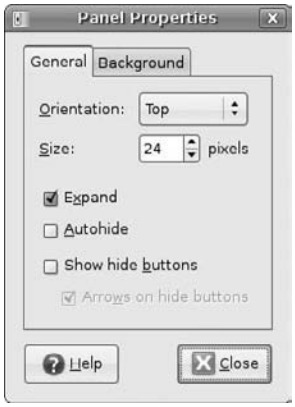
The next two sections explain how to modify the properties of the existing panels on your screen; add other applets, buttons, and controls to those panels; and highlight some of the most fun and useful of these that are provided with Ubuntu Linux.

Customizing Panel Properties

Selecting the Properties entry on any panel's context-sensitive menu displays the dialog shown in Figure 5-19.

FIGURE 5-19

The gnome-panel application's Panel Properties dialog



This dialog contains two tabs: a Background tab, which enables you to control the color and transparency of the current panel or select an existing image to use for its background, and the General tab, shown in Figure 5-19, which enables you to control the location and behavior of that panel. The properties that you can configure on this tab are as follows:

- **Orientation:** Enables you to specify the edge of the screen at which the currently selected panel should display. Selecting another value immediately moves the currently selected panel to the specified edge.

TIP **You can also move an existing panel to another edge of your screen by left-clicking on its background, holding down the left mouse button, and dragging it to the new edge.**

- **Size:** Enables you to specify the height (or width, depending on the panel's orientation) in pixels of the currently selected panel.
- **Expand:** Selecting or deselecting this toggle checkbox determines whether the panel occupies the entire edge of the screen or whether it takes up only that portion of the edge that is required to display the objects that it currently holds.
- **Autohide:** Selecting or deselecting this toggle checkbox determines whether the panel is displayed only when the cursor is moved to the edge of the screen that the currently selected panel is associated with. If this property is selected, the panel is not displayed when the cursor is located over any other portion of the screen.
- **Show hide buttons:** Selecting or deselecting this toggle checkbox determines whether buttons for collapsing or expanding the currently selected panel are displayed at each end of that panel. Selecting this checkbox also activates the Arrows on hide buttons option, which simply determines whether the icons at either end of the panel display arrows to highlight the direction in which the panel will expand or contract if they are selected.

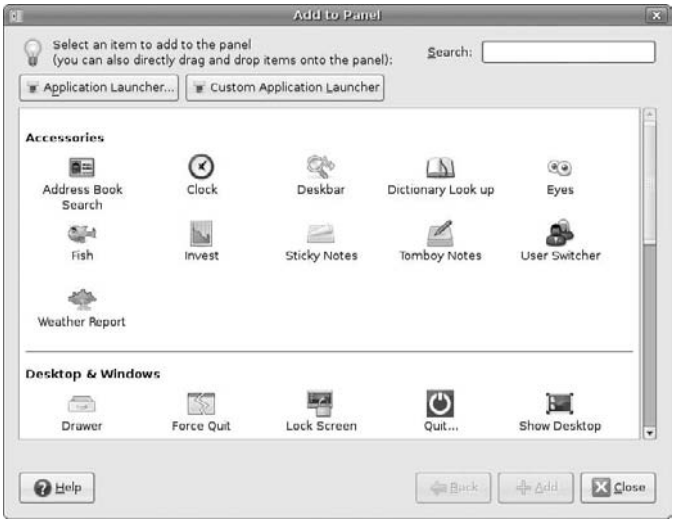
Changes that you make to the properties of any panel take effect immediately. To close the Panel Properties dialog, click Close.

Customizing Panel Contents

Selecting Add to Panel from any panel's context-sensitive menu displays the Add to Panel dialog, as shown in Figure 5-20.

FIGURE 5-20

The Add to Panel dialog



This dialog enables you to add application launchers, buttons, and common applets to any panel that is currently displayed on your screen. All of the currently available buttons and applets are displayed in this scrollable dialog. To add any of these to any panel that is currently visible on the screen, left-click on that applet, hold down the left mouse button, and drag the applet to the location on the panel where you want it to be displayed. Releasing the mouse button positions the selected applet at that location in the panel.

TIP **After adding any applet to any panel, you can move most applets by right-clicking on the applet and selecting Move from its context-sensitive menu. (If the Lock to Panel item is selected on the context-sensitive menu, you will have to deselect this first in order to move the applet. You can then reselect it once you've moved the applet.) After selecting the Move menu item, the applet will follow the current cursor position — once you have moved the applet to its new location, left-click to position it at the new location. Applets can be moved to any other panel that is currently displayed on the screen.**

The applets displayed in the Add to Panel dialog are organized into general categories such as Accessories, applets for Desktop & Windows manipulation, applets for viewing System and Hardware information, and generic Utilities. Some of my personal favorite applets are as follows:

- **Deskbar applet:** Located in the Accessories section, this applet makes it easy to perform a single search of multiple Internet and filesystem locations, query the Tracker desktop search application, query online dictionaries, and even perform specific system actions based on your keywords.
- **Drawer:** Located in the Desktop & Windows section, a drawer is a convenient way of adding multiple applets or buttons to a panel without consuming all available panel real estate. A drawer is essentially just an extension of a panel that can hold other panel objects. Clicking on a drawer causes it to expand, showing the panel objects that it contains, which you can then select just as if they were directly located on a panel. If you're really into applets but don't want to use too many panels, Drawers can be a big convenience.
- **Sticky Notes:** Located in the Accessories section, this applet enables you to create and manage multiple onscreen sticky notes that are analogous to the Post-Its that festoon everyone's monitors and physical desktops nowadays. This applet provides a simple alternative to Ubuntu's Tomboy Notes application, which is located on the Applications Accessories menu.
- **System Monitor:** Located in the System & Hardware section, this applet can be configured to display graphical and associated tool tips that provide information about processor, memory, network, swap space, and hard drive usage, as well as general system load information.
- **Volume Control:** Located in the System & Hardware section, this applet provides an easy way of increasing (or occasionally decreasing) your sound card's volume. (This applet is installed by default in the top panel on Ubuntu systems.)
- **Weather Report:** Located in the Accessories section, this applet displays information about the temperature and general weather conditions in any specified location. Because many of us don't have windows in our offices (or perhaps offices at all), this can be very handy when you're thinking about going outside. It's sad how often I find this applet to be useful — "Oh, it's snowing outside?"

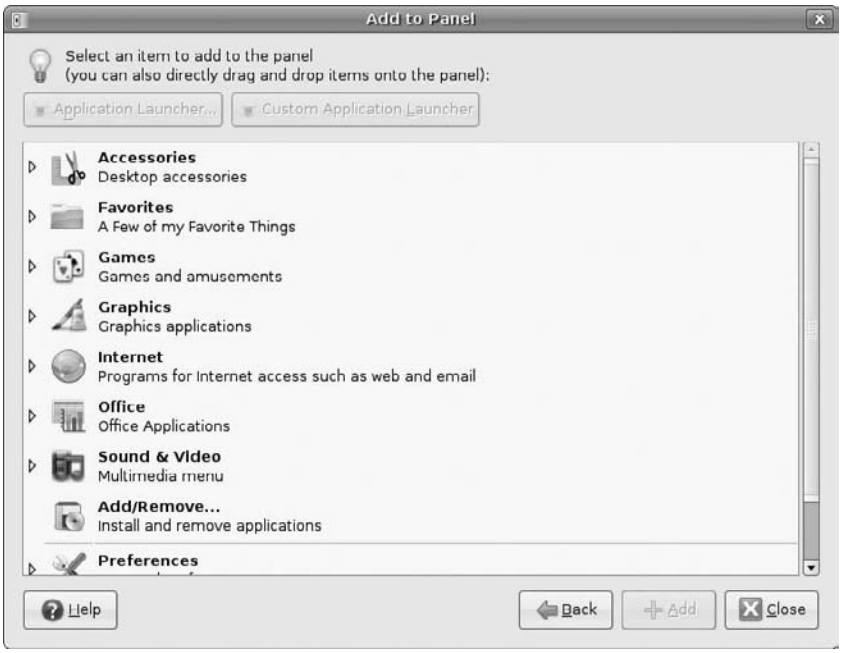
Application launchers are panel entries that execute a specific application for you. Ubuntu comes with several predefined application launchers that you can choose to add a shortcut for starting that application from any panel. To add one of these predefined application launchers, click the Application Launcher button in the Add to Panel dialog. The Application Launcher dialog displays, as shown in Figure 5-21.

To add a predefined application launcher, simply select it from this dialog's scrollable list. To return to the main Add to Panel dialog, click Back.

TIP **Each panel applet has its own configuration opportunities — for example, the Weather Report applet requires that you identify your location to be able to display accurate information about the local weather, but that information isn't really germane to any other applet. To configure any panel applet, right-click on its panel icon and select the Preferences item from the context-sensitive menu that displays. The configuration dialog for that applet displays, showing any configurable properties of that applet.**

FIGURE 5-21

Adding a predefined application launcher to a panel



The Add to Panel application also enables you to add custom application launchers that you can define yourself. This enables you to create panel entries that will launch your favorite applications with your own favorite command-line options. To define a custom application launcher, click the Custom Application Launcher button in the Add to Panel dialog. The Create Launcher dialog displays, as shown in Figure 5-22.

The Create Launcher dialog enables you to set various characteristics of the launcher that you are defining, the application that it starts, and how that launcher displays in the panel. After providing that information, click OK to create the new launcher. By default, the new launcher is displayed in the panel from which you started the Add to Panel application.

Once you're finished customizing and fine-tuning the applets, buttons, and launchers that are displayed in your system's panels, click Close to close the Add to Panel dialog and get back to work.

FIGURE 5-22

Defining a custom application launcher



TIP **To remove an applet from any of your panels, right-click its name and select the Remove from Panel command from its context-sensitive menu.**

Configuring the Screensaver

There was a time when computer monitors were quite sensitive to displaying the same characters for long periods of time. Monitors that showed the same applications in the same locations for long periods of time would end up damaging the phosphors on the inside of the monitor, which therefore displayed a ghostly outline of that application on the screen. (This was known as *screen burn* or *video burn*.) To eliminate or, at least, reduce this problem, many companies developed applications known as screensavers, which would start automatically after a computer system had been unused for a specified period of time, and would display random patterns or images on the screen, preventing the burn-in of default images.

Today's monitors are much more sophisticated than the monitors of old, and are therefore less susceptible to image burn-in. Today, screensavers are typically used because they look nice or for security purposes, so that unauthorized users can't see what is displayed on the screen of your computer while you're at lunch.

By default, your Ubuntu system comes configured to blank the screen after 10 minutes of keyboard and mouse inactivity. You can select a specific screensaver for your system to use (there are hundreds), configure the period of time after which a screensaver is activated, or disable the screensaver on your system entirely. All of these are done through the System Preferences menu's Screensaver Preferences application, which is shown in Figure 5-23.

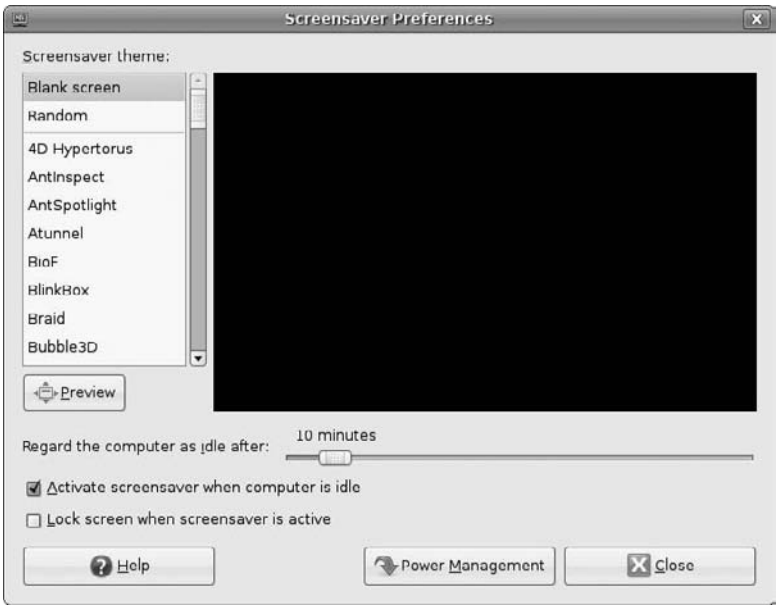
Once this dialog displays, you can explore all of the other screensavers that are available on your system by selecting them from the scrollable Screensaver pane. A preview of the appearance and behavior of the currently selected screensaver displays in the preview pane to the right of the screensaver list.

Once you select the screensaver that you want to use, the Screensaver Preferences dialog also enables you to set some general characteristics of your system's use of that screensaver. To change the period of time after which the screensaver is activated, drag this dialog's slider to reflect the new time period. To deactivate the screensaver entirely, clear the Activate screensaver when session is idle checkbox.

Ordinarily, moving the mouse or pressing any key on your keyboard causes the screensaver to stop displaying its default image and reset its inactivity timer. To combine the screensaver with a dialog that requires that you provide your password to stop displaying the screensaver and regain access to your keyboard and screen, select the Lock screen when screensaver is active checkbox. You will need to supply your password to redisplay your normal desktop.

FIGURE 5-23

Configuring your system's screensaver



The Screensaver Preferences dialog also provides a Power Management button that starts the Power Management Preferences application that you can also execute by selecting the System Preferences Power Management menu item. This dialog enables you to configure how your system reacts to longer periods of inactivity than those covered by the screensaver. Configuring power management is discussed in the next section.

Once you are finished configuring the screensaver, click Close to terminate the Screensaver preferences application.

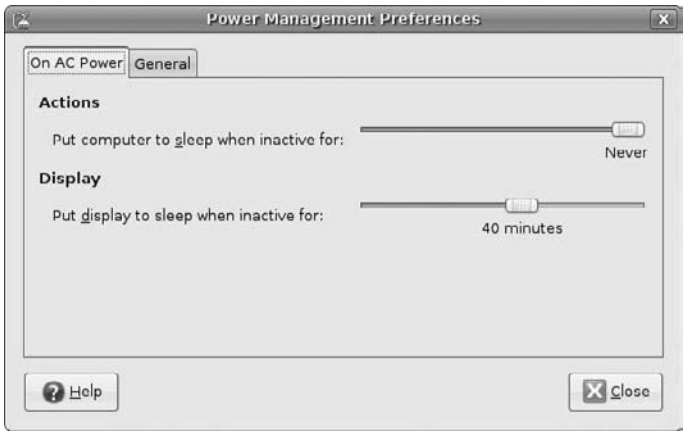
Configuring Power Management

As mentioned in the previous section, the Screensaver Preferences dialog also provides a Power Management button that starts the Power Management Preferences application. A more common (and more intuitive) way of configuring power management and related activities is by selecting the System Preferences Power Management menu item. Selecting either of these displays the dialog shown in Figure 5-24.

The Power Management Preferences dialog provides two or three tabs: the On AC Power tab which enables you to configure how your system behaves after significant periods of inactivity when it is running from AC power (i.e., when it is plugged in), On Battery Power, which is displayed only on laptop systems and which enables you to configure how your system behaves after significant periods of inactivity when it is running from battery power, and the General tab, which enables you to configure how your system should react when the Power or Suspend buttons are pressed, and whether a power management icon is displayed in the notification area at the upper right of the top GNOME panel.

FIGURE 5-24

Configuring power management on your system



The On AC Power tab displays when you first start the Power Management Preferences application, and contains two options:

- **Put computer to sleep when active for:** A slider that enables you to configure your system to go to sleep if it is inactive for the specified period. The default setting for this option is Never, which means that your system will never go to sleep, regardless of how long it has been inactive. If you do not need to access your system remotely (via a network connection or modem), you may want to set this to a slightly lower value — by dragging the slider to the left, you can tell your system to go to sleep after any number of minutes up to one hour. If you need to access your system remotely and will be accessing it over a network, you can still safely set this to a lower value if your system has a BIOS configuration option that tells it to wake up if network activity to the machine is detected. Otherwise, if you need to be able to access your system remotely, you should leave this setting at its default value so that it is always awake and available for access.
- **Put display to sleep when inactive for:** A slider that enables you to configure your system to put your monitor into a low-power mode after it has been inactive for a specified period. To resume regular use (known as *waking* your monitor), you can simply move your mouse or press any key on your keyboard. Putting your monitor to sleep is very different from simply blanking your monitor, as the default GNOME screensaver does. In the latter case, the monitor is simply blank, but is still using a normal amount of power. In sleep mode, the monitor is actually consuming less power. The default setting for this option is 40 minutes, but something like 20 minutes is probably a better selection in most cases.

The On Battery Power tab displays only if you are using a laptop or similar computer, and offers the same settings as the On AC Power tab. In addition, it enables you to select the actions that your system should take when the laptop lid is closed and when battery power is critically low. Your options for these are the same as those provided on the General tab — read on.

The most interesting settings on the General tab are in the Actions area, which defines how your system reacts when you press the Power or Suspend button on your system. (All systems have a Power button; not all systems have a Suspend button.) Your choices for the Power button are Ask Me, which displays a dialog that prompts you to ask what you want to do, and Suspend, Hibernate, and Shutdown. The differences between suspending, hibernating, and shutting down your computer are as follows:

- **Suspend** (also known as *sleep* mode) is a reduced-power mode in which all active applications, associated documents, and the current state of the kernel are saved to RAM, which enables the system to power down most of the devices on your system (such as the hard drive, display, network interfaces, and so on) while waiting for some activity. Your system's processor and memory are still using power while your system is in sleep mode, but these use substantially less power than your system normally does. Be careful when using this mode for an extended period of time on laptop computers. If you are using suspend mode on a laptop without being connected to AC power, you will still lose any unsaved work when your battery is depleted.
- **Hibernate** mode saves all active applications, associated documents, and the current state of your kernel to your hard drive and then actually shuts down the system. When you power the system back on, the system detects the saved state information on your hard drive and restores it to memory, enabling you to pick up exactly where you left off. No power is consumed while the computer is hibernating because it is actually turned off, but restoring the system from the hard disk (known as *resuming* the system) requires that a large file be read and restored to memory, which takes longer than waking up from suspend mode. By default, system images are saved to swap space, so the amount of swap space available on your system should be equivalent to the amount of memory on your system. See the `/etc/acpi/hibernate.sh` script for detailed information about what happens when your system hibernates.
- **Shutdown** terminates all active applications, unmounts all filesystems, and turns your system off. Restarting your system takes longer than waking from suspend mode or resuming from hibernation mode because your system must go through the entire startup process, check your filesystems for consistency, and so on.

You can use the Notification Area preferences in the Power Management Preferences application to configure whether a power management icon is always shown in the Notification Area or whether it is never displayed. Normally, the power management icon is displayed only when your system is running on battery power.

After changing any of these values, click Close to save the new settings and close the Power Management Preferences application.

TIP **Always save any modified files if you know that you are going to be away from your computer for an extended period of time. While the suspend and hibernate modes usually work correctly, you shouldn't take the chance that they do not restore in-progress documents correctly. Your boss, publisher, or thesis advisor probably will not accept "suspend mode ate my homework" as a valid excuse for being late or losing critical data. As a public service, I have personally tested this excuse for you, and it does not work.**

Debugging Power Management Problems

Power management problems can be difficult to troubleshoot because they involve changing the state of your system. In some cases, requesting a suspend or hibernate operation can cause an error that will terminate your current session without any chance to see the error message. If you are experiencing problems getting your system to suspend or hibernate, here are a few things to try:

1. Verify that your kernel actually supports suspend/hibernate. You can do this with the `lshal` command, which queries your system's hardware abstraction layer to identify various capabilities, as in the following example:

   ```
   $ lshal | grep can_suspend
     power_management.can_suspend = true  (bool)
     power_management.can_suspend_hybrid = false  (bool)
   power_management.can_suspend_to_disk = true  (bool)
     power_management.can_suspend_to_ram = true  (bool)
   ```

 This shows that the kernel supports both suspending to RAM (suspend) and suspending to disk (hibernate). These capabilities are enabled by default in official Ubuntu kernels, but if you have built your own kernel, you will want to verify that they are enabled (and enable them and rebuild the kernel if they are not).

2. Verify that the ability to suspend and hibernate is still active in your GNOME environment, as in the following example:

   ```
   $ gconftool-2 -R /apps/gnome-power-manager | grep can
     can_hibernate = true
     can_suspend = true
   ```

 This shows that the `gnome-power-manager` application believes that the system can both suspend and hibernate.

3. If you are only having problems hibernating, make sure that the amount of swap space on your system matches or exceeds the amount of memory in your system. By default, the system image file that is created before your system can hibernate is saved to a contiguous portion of your system's swap space. If you do not have sufficient swap space, you can either increase it (and do so contiguously), or try using another of the hibernation methods that are available for Ubuntu such as those in the hibernate, suspend2, and uswsusp packages. Increasing swap space involves booting without it, resizing the swap and following partitions on your hard drive, reformatting the swap area, and rebooting with the new swap space.
4. Verify that your ACPI settings are correct by examining the file /etc/default/acpi-support. This file sets environment variables used by the ACPI scripts, and may help identify problems. Make sure that both the ACPI_SLEEP and ACPI_HIBERNATE options are set to true. Depending on your hardware, you may need to modify settings such as ACPI_SLEEP_MODE, DOUBLE_CONSOLE_SWITCH, and RESET_DRIVE to get the suspend/hibernate operations to work correctly on your system.
5. If suspend and hibernate both look as though they should work, the next step is to try to run the suspend and/or hibernate scripts manually to see any error message(s) that they display. This is difficult to do on the machine that you are trying to suspend or hibernate, but is easy to do if you

continued

continued

connect to that machine from another via the network using ssh or telnet, and then run the script manually. The scripts that you want to run are those in `/etc/acpi/suspend.d` when suspending your system, and the single script `/etc/acpi/hibernate.sh` when hibernating your system. If you are trying to suspend your system, you would run those scripts from your network connection using a command like the following:

```
$ for SCRIPT in /etc/acpi/suspend.d/*.sh; do \
    sh -x $SCRIPT \
done
```

Using `sh -x` runs all of the scripts in verbose mode to ensure that you see each command that the scripts try to execute, to make it easier to identify problems.

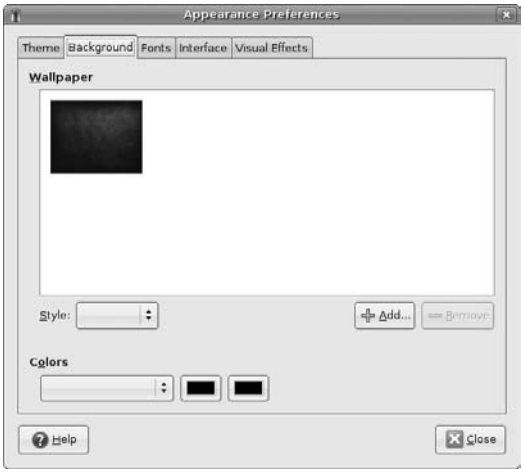
If the machine on which you are working isn't on a local network or you don't have another machine that you can use to connect to that machine, you can try using the script command or redirecting the output of the script into a file to capture any error messages that may be displayed.

Changing Desktop Backgrounds

GNOME also makes it easy for you to change the background displayed on your system's desktops. The background is a solid color or image that is displayed on your screen and over which all other application windows are displayed. To do this, select the System Preferences menu's Appearance menu entry. The Appearance Preferences application displays. To change your desktop background, select the Background tab, as shown in Figure 5-25.

FIGURE 5-25

Configuring your desktop background



Changing the background of your desktop is one of the more common configuration tasks when personalizing your system. To customize your desktop, you can select a solid background color, select an existing graphics image from the default list given in the Desktop Wallpaper dialog, or click Add to navigate to the location of a specific graphics file that you may want to use as your background. (I tend to use vacation or other scenic photos.)

Once you've selected an image, the Style value determines how that image is used as the background. Your options are as follows:

- **Centered:** The selected wallpaper is displayed in the center of your screen at its original resolution.
- **Fill Screen:** The selected wallpaper is increased in size as needed to completely fill the background of your screen. This does not preserve the aspect ratio of your original image.
- **Scaled:** The selected wallpaper is stretched to fit your screen without changing its aspect ratio.
- **Tiled:** The selected wallpaper, at its original size, is repeated multiple times across your screen as needed to fill the background.
- **Zoom:** The selected wallpaper is increased in size so that it completely fills your screen. If the image is not exactly the right size, some portions of the image may be clipped so that the image completely fills the screen.

As you select different wallpaper and style options, the background changes instantly to reflect the current selections. Once you're happy with your background, click Close to close the Appearance application.

Switching Themes

The word "theme" is the term for the combination of window and dialog controls, a color scheme, and a set of icons that give your desktop and the windows that it displays a specific look and feel. Many themes also come with a specific desktop background to complete the visual experience that their author was trying to achieve. Experimenting with different themes can be a lot of fun, and can also be extremely useful to make it easier to read window titles and generally adjust contrast on your screen.

To change or configure the theme used on your Ubuntu system, select the System Preferences menu's Appearance menu item. This application displays the Theme tab when it starts, as shown in Figure 5-26.

Ubuntu's GNOME desktop comes with many attractive themes. By default, Ubuntu uses a theme that is designed to be attractive, easy on the eyes, and easy to use. To switch to another of the themes provided with Ubuntu Linux, select it from the scrollable theme pane in the center of the dialog. Even though the theme pane includes preview images of each available theme, the new theme takes effect immediately so that you can see what it will actually look like. Figure 5-27, for example, shows the Theme tab after selecting the Crux theme. You'll have to trust me that the window decoration background is now purple, but you should be able to see the differences in the window decoration, font placement, highlighting, window controls, dialog button shapes, and so on.

Once you've selected the theme that you like best at the moment, click Close to close the Appearance Preferences dialog. Your new theme will be used each time you log in, until you change it.

FIGURE 5-26

Specifying theme preferences

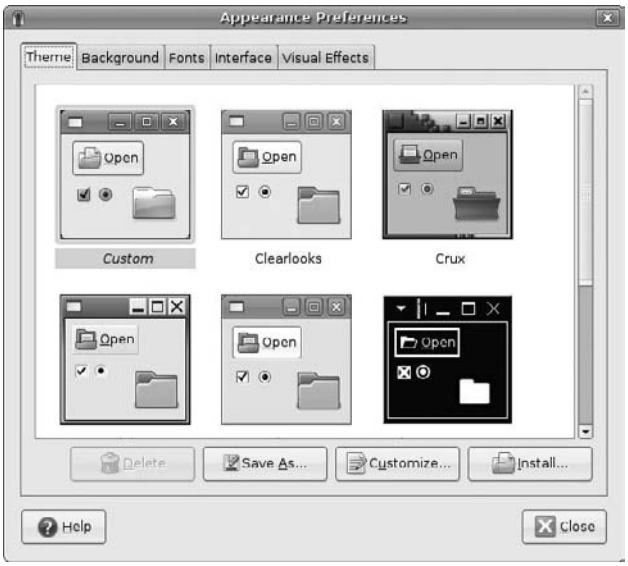
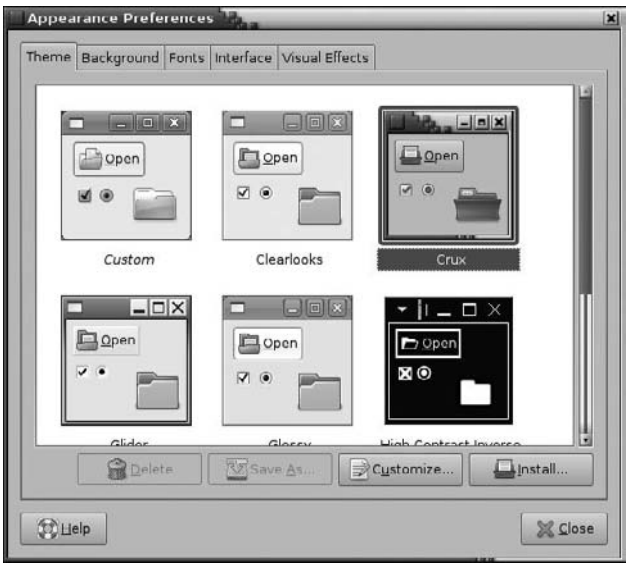


FIGURE 5-27

Selecting a different window theme



TIP **Although Ubuntu provides a large collection of nice themes, artistically inclined users are always creating new ones. Many excellent (or just plain interesting) themes are available at a variety of Web sites. My favorites are the GNOME Art site (`http://art.gnome.org/`), GNOME-Look.org (`www.gnome-look.org/`), and Freshmeat's GNOME themes section (`http://themes.freshmeat.net/browse/58/`). A quick Web search for "GNOME themes" with your favorite search engine will turn up many more sites. To add any of these themes, simply select the Install Theme button after downloading the theme and navigate to its location on your system. The new theme is added to the scrollable themes list, and can be selected like any other theme. Somewhere out there is the perfect theme for you!**

Assistive Technologies for Using GNOME

If you have trouble seeing the screen or using a keyboard or mouse, the GNOME and Ubuntu folks are sensitive to your needs and offer various assistive technologies that do things such as providing onscreen keyboard support, and supporting screen magnification and audio output from portions of the screen. These features are installed by default as part of all Ubuntu desktop installations, but are not activated by default.

You can activate the capabilities of these packages by selecting the System Preferences Universal Access menu item, which displays the dialog shown in Figure 5-28. You can then enable selected items, and log out and log back in to activate them. A separate Keyboard Accessibility item on the Universal Access dialog provides a shortcut to the same dialog that is displayed if you click the Keyboard Accessibility button shown in Figure 5-28.

FIGURE 5-28

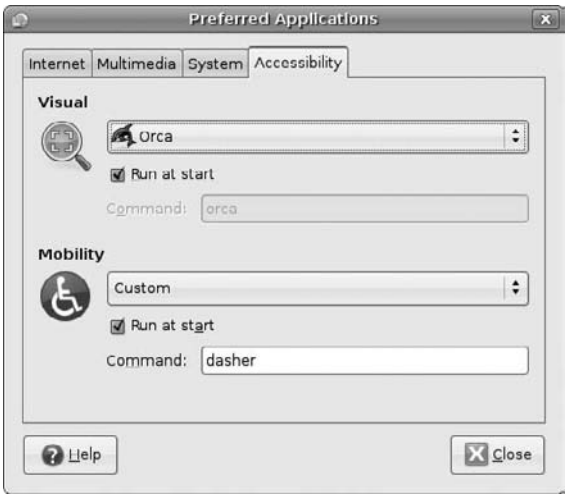
Enabling Assistive Technologies in GNOME



Clicking the Preferred Applications button displays a dialog, shown in Figure 5-29, that enables you to define the applications to invoke in response to certain Internet, multimedia, and system events, and which also provides an Accessibility tab that enables you to activate features such as screen reading and magnification via the Orca application (`http://live.gnome.org/Orca`) and a rapid text entry system known as Dasher (`www.inference.phy.cam.ac.uk/dasher/`) to help simplify text-entry when a standard keyboard cannot be used easily.

FIGURE 5-29

Customizing accessibility in GNOME



These applications aren't perfect, but are actively under development to guarantee that anyone can take advantage of the power of Ubuntu Linux. For more information about assistive technologies and GNOME, see `http://projects.gnome.org/accessibility/`.

GNOME Keyboard Shortcuts

If you have problems controlling a mouse or simply prefer to keep your hands on the keyboard at all times, GNOME provides several special key sequences that enable you to perform most traditionally mouse-driven desktop activities using only the keyboard. These are commonly referred to as *keyboard shortcuts*. To activate these, press and hold down the first key in the sequence while pressing the second. The most popular keyboard shortcuts for the GNOME desktop are as follows:

- **Alt+F1:** Displays the Applications menu.
- **Alt+F2:** Displays the Run Applications dialog, in which you can enter the name of a specific application that you want to execute.
- **Alt+F4:** Closes the window that is currently selected.
- **Alt+F5:** Reduces the size of the current window if it is maximized.
- **Alt+F7:** Enables you to move the window that is currently selected. After pressing this key combination, use the arrow keys to move the window to a new location and press any other key on the keyboard to complete the move operation.
- **Alt+F8:** Enables you to resize the window that is currently selected. After pressing this key combination, use the arrow keys to increase or decrease the size of the window and press any other key on the keyboard to complete the resize operation.
- **Alt+F9:** Minimizes the window that is currently selected.

- **Alt+F10:** Maximizes the window that is currently selected.
- **Alt+Tab:** Cycles through all of the windows and dialogs that are currently displayed on the screen, highlighting each in turn (known as *giving the focus* to that window) until you find the one that you're interested in.
- **Alt+Spacebar:** Displays any context-sensitive menu associated with the window that is currently selected.
- **Ctrl+L:** When browsing for a file in any File Open dialog, displays a pop-up box in which you can type the full pathname of the directory containing the files that you are looking for, saving you many clicks if you need to drill down to a deeply nested directory. This dialog does automatic directory-name completion as you type, and is one of my favorite GNOME features.

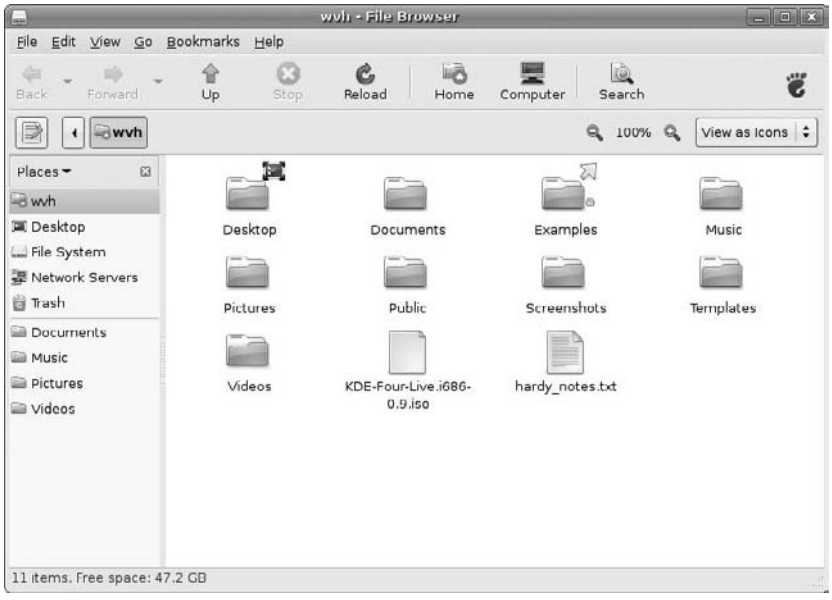
GNOME supports several other keyboard shortcuts, most of which perform standard actions within various applications. For a complete list of available keyboard shortcuts, see the GNOME Accessibility Guide, available by selecting the System Help System Documentation menu item.

Introducing the Nautilus File Manager

The Nautilus file manager, known to its friends simply as Nautilus, is a convenient application for computer users who are used to graphically displaying, examining, and moving around in files and directories on their computer systems. Nautilus is GNOME's equivalent of the Microsoft Windows Explorer or the Mac OS X Finder, and makes it similarly easy to work with and modify your files and directories. Figure 5-30 shows Nautilus exploring the contents of my home directory on one of my systems.

FIGURE 5-30

Exploring my home directory in Nautilus



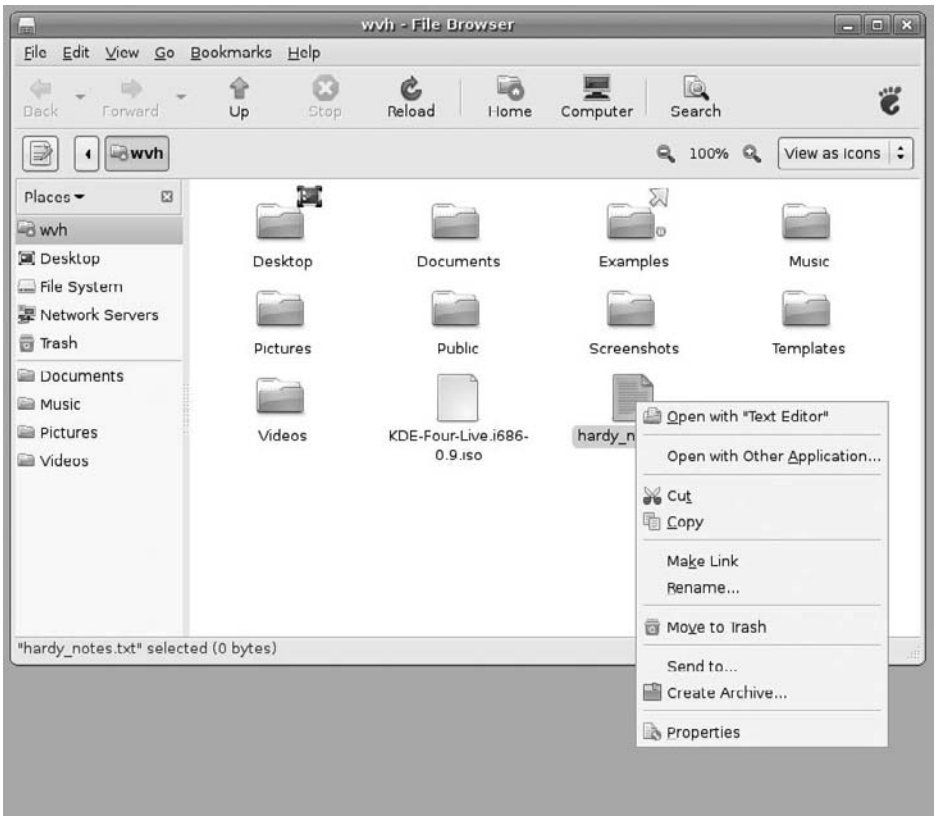
As mentioned earlier in this chapter, the Places menu in the top panel on the default GNOME desktop shipped with Ubuntu systems is actually a collection of shortcuts to using Nautilus to explore different portions of your system or network. For example, I displayed my home directory in Nautilus by selecting Home Folder from the top panel's Places menu.

Basic Operations in Nautilus

Like other graphical file managers, Nautilus makes it easy to manage and manipulate files and directories. Most of the commands related to working with files and directories are located on the Edit menu, and are also accessible from the context-sensitive menu displayed when you right-click on any object in the Nautilus display pane. Figure 5-31 shows the context-sensitive menu displayed when right-clicking on a file with the .txt extension in my home directory.

FIGURE 5-31

A context-sensitive menu in Nautilus



As you can see, this menu provides the standard actions that you'd expect to find in any file manager, enabling you to do the following:

- Open the file with the default application associated with that type of file.
- Open the file with other applications associated with that type of file.
- Cut or copy the file, enabling you to paste the file (or a copy of it) into another location to move or copy the file.
- Create a symbolic link to or rename the selected file.
- Rename the file.
- Delete the file by moving it to the Trash Can, which is a temporary repository for files before they are actually purged from your system.
- Send the file to a specified device or create an archive file in compressed tar format containing that file. (See Bonus Chapter 6, "Backing Up and Restoring Files," for information about Linux archive file formats.)
- Examine or modify the properties of the file, including its name, permissions, and how it is displayed in Nautilus.

Examples of Using Nautilus

The previous section described the most common operations that you can do in Nautilus; this section provides a few examples of such operations to get you started working with Nautilus if you haven't used graphical file managers before.

Exploring the Filesystem

The GNOME desktop's Places menu contains several preconfigured shortcuts to various locations on your system and on your local network. If you're lucky, the directory or network location that you want to explore will be listed there. If not, you can navigate to any location on your system in Nautilus by doing the following:

1. Open Nautilus in a predefined location by selecting it from the Places menu.
2. Drill down into subdirectories of the selected location and double-click on the icon for that folder.
3. Move up a level to the directory that contains the directory you are currently viewing and click the Up button in the Nautilus toolbar.

Moving around between different files and directories in Nautilus is easy once you get the hang of it, and will soon become second nature to you.



You can start Nautilus and instantly display the contents of any accessible directory by starting Nautilus from the command line and specifying the name of that directory as an argument. For more information about using the command line, see Chapter 14, "Using Command-Line Tools."

Copying Files or Directories

To copy a file or directory in Nautilus, do the following:

1. Navigate to the directory that contains the file or directory that you want to copy.
2. Left-click on the icon for that file or directory to select it.

3. Select the Copy command from the Edit menu or right-click on the selected icon and select Copy from the context-sensitive menu.
4. Navigate to the directory where you want to place the copy of the selected file or directory.
5. Select the Paste command from the Edit menu or right-click on the background of the current directory in Nautilus and select Paste from the context-sensitive menu.

At this point, the icons for the files and directories that you had selected for copying display in the current Nautilus window, and the copy operation is complete. Copying directories that contain many files may take a bit of time.

Don't worry about accidentally overwriting files with the same name — if the directory to which you are copying files already contains files or directories with the same names as the files or directories that you are copying, the duplicate file will have the word (copy) added to its name to uniquely identify it.

Selecting Multiple Files and Directories

As you work with Nautilus, you may find that you want to move or copy multiple files and directories at the same time. This is easy enough to do as long as you select all of the icons associated with those files and directories before performing a copy or cut operation. You can select multiple items in Nautilus in a variety of different ways as follows:

- To select multiple files and directory icons, hold down the Ctrl key on your keyboard while selecting each of the icons for those files and directories.
- To select all of the files and directories that are alphabetically between two files and directories, select one icon, hold down either Shift key, and select the icon at the end of the range of files that you want to select.
- To graphically select multiple files that are next to each other in the Nautilus display windows, left-click the Nautilus background beside any of those files, hold down the left mouse button, and drag the mouse cursor to display a rectangle that contains all of the files that you want to select.

You can also combine these selection operations. For example, you could select a range of files graphically by using the Shift key and then holding down the Ctrl key to add other files or directory icons to the current selection.

Moving Files or Directories

To move a file or Directory in Nautilus, do the following:

1. Navigate to the directory that contains the file or directory that you want to move.
2. Left-click on the icon for that file or directory to select it.
3. Select the Cut command from the Edit menu or right-click on the selected icon and select Cut from the context-sensitive menu. The icons for the selected file or directory are removed from the current Nautilus window.
4. Navigate to the directory where you want to place the selected file or directory.
5. Select the Paste command from the Edit menu or right-click on the background of the current directory in Nautilus and select Paste from the context-sensitive menu.

At this point, the icons for the files and directories that you removed from the first directory will display in the current Nautilus window, and the copy operation is complete. Copying directories that contain many files may take a bit of time.

Renaming a File or Directory

To rename a file or directory in Nautilus, do the following:

1. Navigate to the directory that contains the file or directory that you want to rename.
2. Right-click on the selected icon and select Rename from the context-sensitive menu. The name of the file will be selected in Nautilus.
3. Enter the new name of the selected file or directory and press Return. The updated name displays in Nautilus.

Creating a Directory

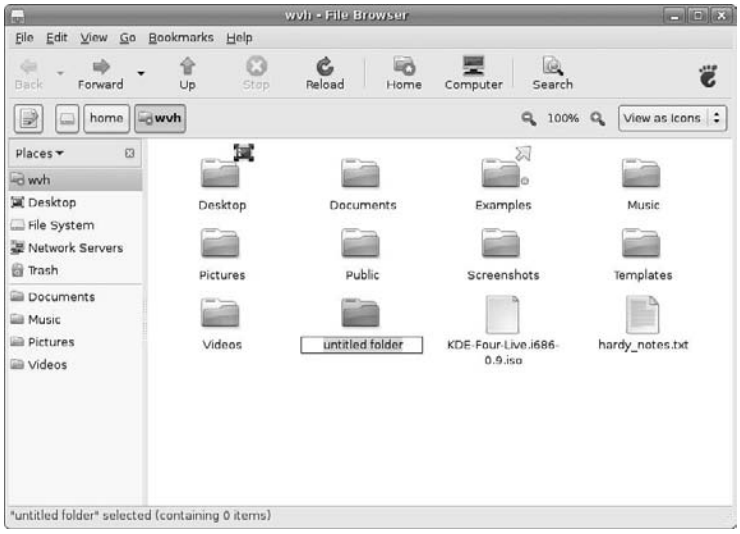
Creating a file is usually done by starting the application associated with that type of file and saving your work with the name of the file that you want to create. Creating directories can either be done from the command line, as explained in Chapter 14, or quickly and easily within Nautilus.

To create a directory in Nautilus, do the following:

1. Navigate to the directory where you want to create the new directory.
2. Right-click on the background of that directory as shown in Nautilus and select Create Folder from the context-sensitive menu, or select File New Folder from the Nautilus menus. A new folder with the default name "untitled folder" displays in the Nautilus window, as shown in Figure 5-32.

FIGURE 5-32

Creating a new folder in Nautilus



3. Note that the name of the new folder is highlighted when that folder is created. To give the new folder a specific name, you can simply type that name or rename the directory as explained in the previous section.

Getting More Information About Nautilus

Nautilus includes extensive and excellent documentation that is available by selecting Contents from its Help menu. Because it is so similar to the types of file managers used on other systems, you should find Nautilus extremely easy to use and master.

Using the Tracker Desktop Search Application

As disks and other local storage media get larger and larger, the amount of data that people save and keep around for an extended period of time gets equally large. Although this anecdote will date me, the first hard drive that I ever bought was a whopping 30MB and I remember thinking "That should last a lifetime! How can I ever fill that up?" It actually didn't take very long to fill the drive, and I remember having to clean up after myself frequently, writing files that I didn't really need to backup media so that I could free up space on my hard drive. Nowadays, some of my data files are larger than that entire drive, and the fact that my desktop system has access to several terabytes of storage means that I'm less likely to clean up after myself anytime soon. It's simply easier to keep everything online (with offsite backups of truly critical data, of course). This is convenient when I want to consult an article about a defunct Linux distribution that I wrote a decade ago, but is not so convenient when I remember writing something a long time ago, but have no idea what file or directory it's in.

Just as search engines make it possible to find information on specific topics regardless of where the information is located on the Internet, desktop search applications make it easy to find information that is located somewhere on your desktop computer. Online search engines and desktop search applications actually work in a very similar way — they crawl the information that they are responsible for and create an index of that information against which you can submit queries.

GNOME-based Linux distributions have provided a desktop search application called Beagle (`http://beagle-project.org/Main_Page`), a desktop search application that was written in C# and which uses Novell's Mono project to enable it to run on Linux systems. Although Beagle can index many different types of files and is nicely designed for integration, in my experience it is extremely CPU-intensive while indexing files — so much so that it is the first process I would look for to terminate if one of my systems were to grind to a halt.

I was not the only person to notice or believe this. Although newer versions of Beagle (and Mono) are much improved and continue to improve, why wait? Ubuntu now uses a desktop search application called Tracker (`www.tracker-project.org/`). Tracker can index text, HTML, PDF, Microsoft Office, and Open Office documents, and can index a variety of audio and video files based on their MIME type.

NOTE **MIME (Multipurpose Internet Mail Extensions) is an Internet standard for identifying file formats, and is associated with mail systems because the need to identify files types was most commonly recognized when people began to attach files to their e-mail messages.**

In addition to its speed and search capabilities, Tracker provides some cool features such as the ability to tag files so that you can later find them by searching for the tag, and provides some automatic capabilities to expand your searches using the root(s) of your search term(s) (known as *stemming*), increasing the chance of finding what you're looking for.

NOTE **For similar information about desktop search applications on Kubuntu systems, see the section "Using the Strigi Desktop Search Application" in Chapter 6. The section explains how to run and use a well-known open source application called Strigi (`http://strigi.sourceforge.net/`), which is used on Kubuntu systems to index your local storage media and to search that index. A GNOME panel plug-in for Strigi is available as a standard DEB package (`http://packages.debian.org/unstable/gnome/deskbar-plugins-strigi`).**

Tracker is written in C, and consists of a daemon that indexes your files, a graphical configuration application, and several command-line and graphical utilities that make it easy to query and monitor the Tracker daemon and its index. The easiest way to query Tracker is by selecting the Magnifying Glass icon in the upper GNOME panel. Tracker is also integrated into the GNOME Deskbar applet (mentioned earlier in this chapter in the section "Customizing Panel Contents," and discussed in more detail later in this section), and can also be queried from the command line, which makes it handy for use in scripts.

Searching with Tracker

The Tracker daemon, trackerd, is automatically started as part of each Ubuntu user's GNOME login session. By default, Tracker automatically searches and indexes all of the files in your home directory. Once Tracker completes its initial scan of your files and directories, you can search the Tracker index in one of four basic ways:

- From the Tracker panel applet, which starts the Tracker Search Tool.
- From the tracker-search-tool application, which you can start from the command line, or by selecting the Accessories menu's Tracker Search Tool menu option, or by creating a desktop launcher for this application.
- From the command line by using the `tracker-search` command. This tool is part of the tracker-utils package, which is not installed by default.
- As part of the deskbar-applet panel tool.

The Deskbar panel applet is the easiest way to submit a query to Tracker, and also provides easy access to the Tracker at any time. (It is also discussed in more detail in the next section.) After adding this applet to a panel (as described in "Customizing Panels," earlier in this chapter), you can click on its icon in the top GNOME panel drops down a query box into which you can type the term(s) that you're looking for, as shown in Figure 5-33.

Pressing Return then automatically starts the graphical Tracker Search Tool application with these as the search terms, as shown in Figure 5-34.

As you can see from Figure 5-34, a Tracker search locates items whose file or path names match the phrase that you're searching for. Double-clicking any of the search results shown in the Tracker Search Tool window will open that result in the application with which files of that type are associated.

The summary line above the list of matching items shown in Figure 5-34 lists 95 available matches for the search term "Ubuntu." If you have a large number of search results, the icons in the Categories section of the Tracker Search Tools can help you determine where various results are located, helping you find the "right" search result. Default categories include: All Files, Folders (a view of the folders in which search results were found), Documents (PDF, HTML, TeX, text, and related files), Images (image files in various formats that match your query), Music (audio files in various formats that match your query), Text (text files only), and Development (source code in various languages, Makefiles, and so on).

FIGURE 5-33

Initiating a search in the Tracker panel applet

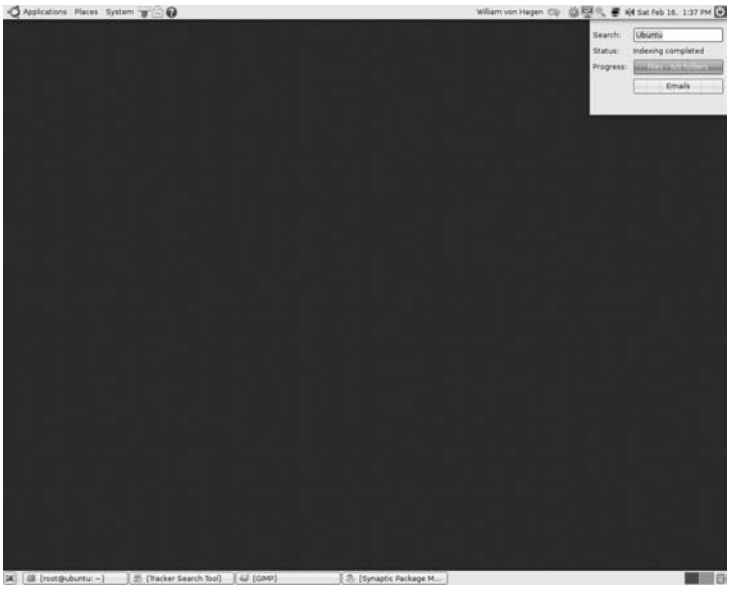
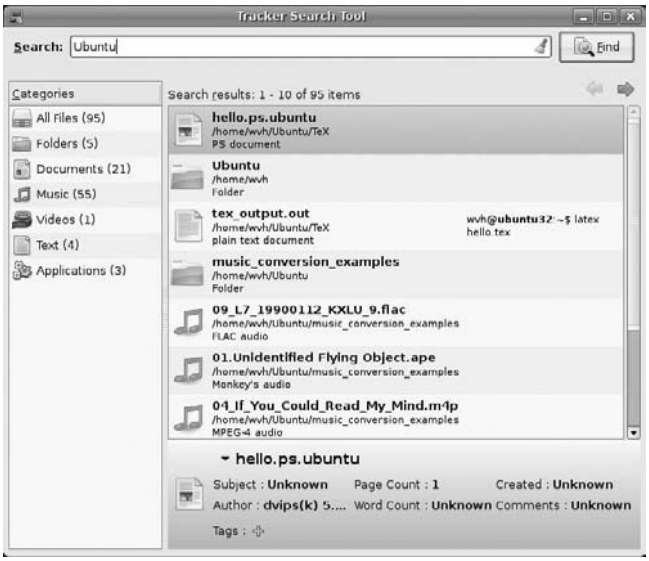


FIGURE 5-34

Search results in the Tracker Search Tool



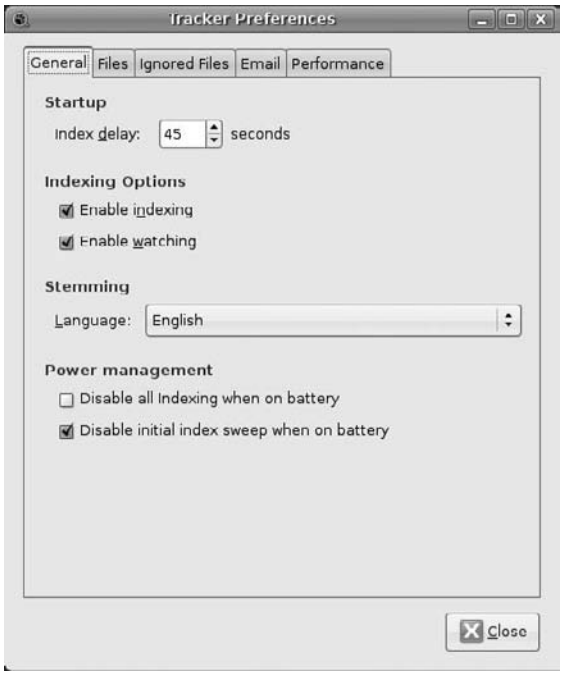
Once started as part of your GNOME login session, the Tracker daemon automatically runs in the background, continuously checking for new files in the directories that it is monitoring, indexing new files as they appear, and updating the index entries for existing files when you modify them.

Configuring Tracker

Tracker is a very powerful and flexible desktop search application whose behavior can be configured using the Tracker Preferences application. Per-user configuration information for Tracker is stored in the file ~/.config/tracker/tracker.cfg, which can be edited (carefully) using a text editor, or automatically by using the Tracker Preferences application. You can start the Tracker Preferences application from the command line (by executing the tracker-preferences command) or by selecting the System➪Preferences➪Search & Indexing menu item. Figure 5-35 shows the initial dialog that displays when you start the application.

FIGURE 5-35

Configuring the Tracker Search Tool



The different tabs in the Tracker Preferences application enable you to configure different aspects of how the Tracker daemon interacts with your system. These tabs are as follows:

- **General:** Enables/disables indexing and whether trackerd actively monitors changes to indexed files and directories. Enables you to set the delay after which indexing begins once you log in. Enables you to set the language used when stemming, which is a mechanism for locating related words by identifying common roots. Enables you to configure how the Tracker daemon behaves on laptops that are running on battery power.

- **Files:** Lets you configure the directories that are indexed by the Tracker. This consists of two sets of directories — those that are indexed by the Tracker daemon and which are actively monitored for changes, and those that are simply indexed once and are not subsequently checked for changes. The settings on this tab also enable you to configure whether the Tracker daemon indexes file contents in addition to filenames (both of these are enabled by default), whether the Tracker daemon generates a thumbnail image for each file that it indexes (disabled by default), and whether the Tracker daemon will cross mount points to index directories that are mounted under a directory that it is indexing or watching.
- **Ignored Files:** Enables you to identify directories and wildcard patterns for filenames that it will not index. For example, you can disable indexing anything under your music directory by entering **/home/***username***/music** in the Ignored Paths text area (you would, of course, replace ***username*** with your login name), and you can disable indexing all MP3 files by entering ***.mp3** in the Ignored File Patterns text box. To add a new path or file pattern to ignore, click Add beside the appropriate text area, enter the path of file pattern, and press Return/Enter. To remove an existing one, highlight it in the appropriate text area and click Remove.
- **Email:** Enables/Disables indexing of the directories used by the Evolution, Thunderbird, and KMail e-mail applications. Only the checkboxes for e-mail applications that are actually installed on your system are selectable on this dialog.
- **Performance:** Enables you to configure how the Tracker daemon operates on your system in order to minimize its impact on the system. Options include Throttling, (which lets you reduce the amount of CPU that the Tracker can use), Resource Usage (the amount of memory that the Tracker daemon uses — reduced memory use causes slower indexing but leaves more memory available to other applications), and Indexing Limits (per file) (how the Tracker daemon handles merging existing indexes with index entries for new files, and the maximum amount of text and number of words that should be indexed in each file).

The most commonly used of these tabs are the Files and Ignored Files tabs. The Files tab is especially useful if you store information that you want to index in directories other than your home directory. The Ignored Files tab is especially useful if you do not want to index specific types of files or files in certain directories. For example, I use a network-attached storage (NAS) device at home to archive old files that I want to be able to access from multiple systems, which I can do by mounting its filesystem on any given machine. Aside from archived personal and research data, I also store my music collection on this system. I therefore add the mount point for this device to the Files tab's list of directories that I want to index and watch, and add the music subdirectory of this mount point to the Ignored Files tab's list of Ignored Paths. I could also do the latter by adding the file extensions for the file formats in which I store music (*.ogg, *.flac, *.shn, and *.mp3) to the Ignored Files tab's Ignored File Patterns list, but this would prevent me from locating any music files that are in any other directories (such as on a specific system), which I personally don't want to do.

As you can see, the configuration options provided by the Tracker Preferences tool give you a lot of control over Tracker's performance and behavior. All in all, I find Tracker to be the best of the desktop search applications. However, one of the reasons for this is that Tracker is nicely complemented by other search mechanisms that are provided as part of the GNOME desktop in the GNOME Deskbar applet, which is described in the next section.

TIP **In addition to the Tracker Applet, Preferences, Search Tool, and command-line tracker-search applications discussed in this section, you can access several other Tracker-related applications by installing the tracker-utils package. This installs a number of command-line utilities, such as the tracker-stats and tracker-status utilities that provide information about the searches that Tracker has performed, how Tracker is performing in general, and so on.**

Using the GNOME Deskbar Applet

Deskbar applet, how do I love thee? Let me count the ways. I love thee when I forget the name of the file in which I've stored the great American novel or my mother's recipe for dilly beans, or when I need to look up a word in a dictionary or thesaurus, or when I want to search the Web. And most of all, I love thee when I don't want to have to start specific applications for each of these different, but related, tasks.

The Deskbar applet is a one-stop shopping applet for any search task that I can think of, including searching for non-file-related data such as specific applications and windows. The Deskbar applet even includes some useful hooks for performing common system tasks, such as logging out, suspending your system, and even shutting it down.

The Deskbar applet is installed by default as part of the top panel on an Ubuntu system, and is accessed by clicking on its "magnifying glass and more" icon. Figure 5-36 shows the Deskbar Applet dialog, which displays after clicking on this icon. (This looks slightly different on different Ubuntu releases, but the functionality is the same.)

FIGURE 5-36

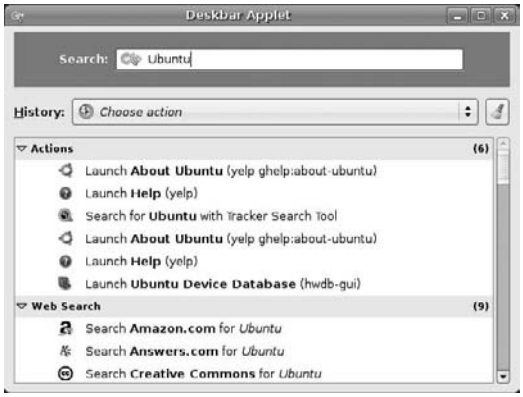
The initial Deskbar Applet dialog



Figure 5-37 shows the default results of a search for the word "Ubuntu" on one of my systems. As you can see from the position of the elevator button in the scroll bar at the right of this dialog, these results are only the tip of the iceberg — there are many more possible results.

FIGURE 5-37

A search for Ubuntu in the Deskbar Applet



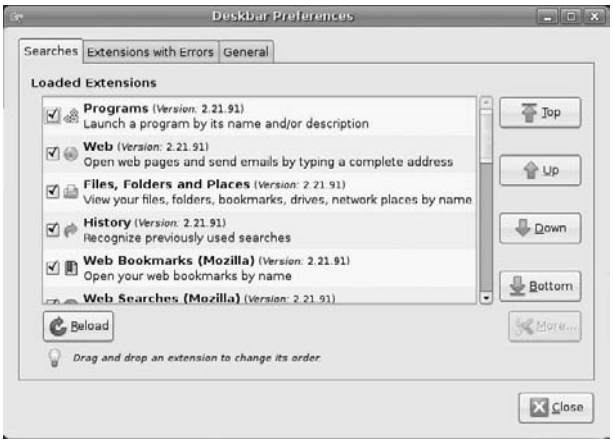
Search results in the default configuration of the Deskbar applet are organized into the following sections (listed alphabetically, not in the order that they are displayed by default):

- **Actions:** Execute various programs in response to matches for your search terms.
- **Documents:** Explore documents whose pathnames or content match your search terms.
- **Files:** Explore files with unrecognizable extensions or MIME types whose path or filenames match your search terms.
- **Music:** Explore files whose MIME type indicates that they are audio data and whose path or filenames match your search terms.
- **People:** Send e-mail to an e-mail address or search various Web sites for that address.
- **Places:** Explore matching items in various directories on your system.
- **Videos:** Explore files whose MIME type indicates that they are video data and whose path or filenames match your search terms.
- **Web:** Open a specific Web site whose URL is one of your search terms.
- **Web Search:** Extend your search to various Web sites.

Simply pressing return will execute the first item in the list, or you can click on a specific result to perform the associated action, whether it's opening a file or performing a Web search for your search terms. Empty sections are not displayed — for example, if no video file matches your search terms, there won't be a Videos section. However, as you can see, there's a whole lotta searching going on, and all of it available through this single applet. Personally, that's too much data for me — I prefer to reduce the overwhelming number of possible locations that are offered in response to a Deskbar applet search by configuring it to search only some of my favorites. To configure the locations that the Deskbar applet will query (or offer as additional query locations), right-click on its icon in the panel and select the Preferences menu option to display the Deskbar Preferences configuration dialog shown in Figure 5-38.

FIGURE 5-38

Configuring search locations for the Deskbar Applet



This dialog lets you configure both the locations that will be queried for your search terms and the order in which those locations will be queried. The different search locations that are available in the Deskbar Preferences dialog are as follows (listed in alphabetical order, not in the order that they are displayed by default):

- **Computer Actions:** Enables you to perform various system actions specified by keyword. Supported keywords include logoff, restart, shutdown, and suspend.
- **del.icio.us Bookmarks:** Searches bookmarks that you have saved on the del.icio.us social bookmarking site.
- **Dictionary:** Looks up search terms using the GNOME Dictionary applet.
- **Files and Folders Search:** Uses your search terms as a pattern to find files and directories whose names match that pattern.
- **Files, Folders, and Places:** Enables you to search your bookmarks, files, directories, mounted drives, and network places by name. This search option is enabled by default.
- **History:** Remembers previous searches that you've made and automatically resurrects your search results. Previous searches are also available from the Choose Action drop-down list shown in Figures 5-36 and 5-37. This search option is enabled by default.
- **Programs:** Enables you to launch a specific application by name or the description that is associated with it on the GNOME desktop. This search option is enabled by default.
- **Recent Documents:** Looks for matches for your search terms in the files listed on the Places➪Recent Documents menu.
- **Templates:** Enables you to create new files based on the names of existing document templates on your system. These templates are the same ones listed on the Create Document menu that is available from the desktop's context-sensitive menu.
- **Tomboy Notes:** Searches any saved notes that you have created with the Tomboy desktop notes application (Applications Accessories Tomboy Notes).
- **Tracker Live Search:** Looks up your search terms as you type, character by character, in the Tracker Search Tool. This search option is enabled by default.
- **Tracker Search:** Looks up your search terms in the Tracker Search Tool. This search option is enabled by default.
- **Web:** Enables you to open Web pages by typing a URL or to send e-mail by typing an e-mail address. This search option is enabled by default.
- **Web Bookmarks:** Enables you to open Web bookmarks by name. This search option is enabled by default.
- **Web History:** Searches the names of any Web sites that are present in your Web browser history.
- **Web Searches (Mozilla):** Enables you to perform Web searches using your browser's search settings.
- **Window Switcher:** Enables you to switch to a different desktop window by typing its name as your search term.
- **Yahoo! Search:** Incrementally searches for your search terms, character by character, using the Yahoo search engine.

Search options that are enabled are preceded by a check mark, and they are searched in the order that they appear. To enable or disable a search option, click on the checkbox. You can also change the order in which search options appear by selecting the name of a search option and dragging or dropping it to a new position in the list, or by using the Top/Up/Down/Bottom buttons to change its position.

Everyone's idea of what and where to search are different, but I typically reconfigure the Deskbar applet to use the following search options, in the following order:

- **Tracker Search:** I'm often looking for data in a file that is somewhere on my system, so putting this option first is a good complement to my Web search options.
- **Dictionary:** Believe it or not, I am occasionally at a loss for words (or at least the right one), so it's useful to be able to look up words in an online dictionary.
- **Web Searches (Mozilla):** Offers to submit my search terms to various search engines and popular sites. I put this third because it offers more possible actions than the first two.
- **History:** Makes it easy to reuse previous searches, in case I want to pursue some subject in more detail, find a specific file again, and so on.

As you can see, the Deskbar applet is a great complement to the Tracker family of desktop search applications. It's extremely convenient to be able to initiate a search for something from a single location, and it's sufficiently configurable to enable you to tailor your search preferences to the locations and types of searches that you typically want to perform.

Using a Window Manager

The first section of this chapter, "What's a Desktop? Graphical Environments for Linux," discussed the differences between window managers and desktops, and hinted at the fact that more X Window system window managers are available than you could shake a stick at. Popular window managers include After Step, Black Box, Compiz, Enlightenment, Fluxbox, Fvwm, ICEwm, kwm, Metacity, twm, vtwm, and Window Maker, but there are many, many more.

After you experiment with GNOME for a while, you may be curious about what these window managers look like and how they work. Half the fun of a flexible computing environment such as Linux and the X11 Window system is experimenting with new software "because it's there." You might be especially interested in window managers if you are running Ubuntu on a slow system or a system with a limited amount of memory, such as many older laptops, because window managers require far less memory and other system resources than complete desktop environments such a GNOME and KDE do.

As you might hope, it is extremely easy to temporarily switch between using GNOME and a specific window manager. First, you must enable the Universe repository, as explained in Chapter 23, "Adding, Removing, and Updating Software," because that is the place where most of the window manager packages are located. Next, you have to install one or more window managers using any of the software update tools that are also discussed in Chapter 23. After installing window managers and logging out, you can start any installed window manager by clicking the Options button at the bottom left of the login window and selecting Select Session from the menu that displays. Figure 5-39 shows a sample session menu on an Ubuntu system where a few window managers (Fvwm, IceWM, and Fluxbox) have been installed.

Select the radio button beside the window manager that you want to experiment with, and click Change Session to close the menu. When you log in, you will see a dialog like the one shown in Figure 5-40, asking if you want to use this window manager permanently.

FIGURE 5-39

The GNOME login screen's Sessions menu

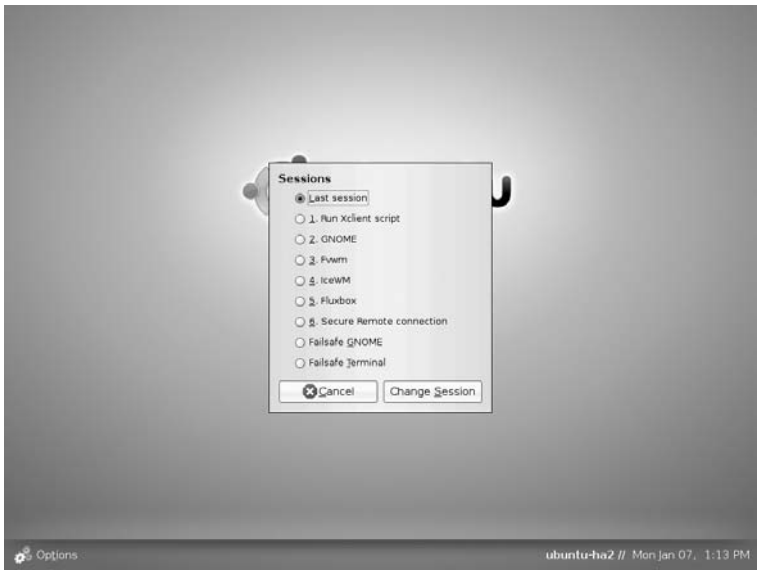


FIGURE 5-40

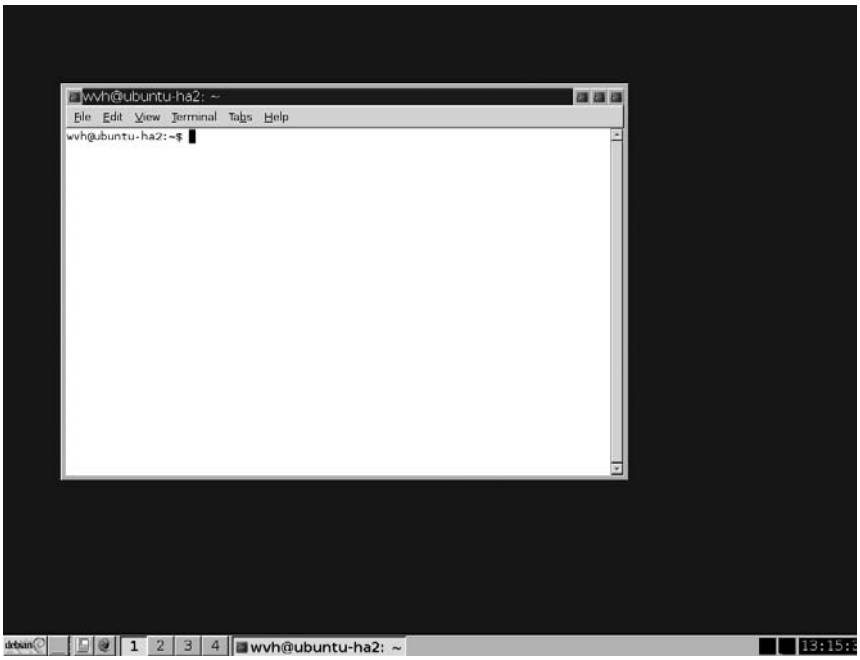
Permanently changing your window manager in the GNOME login screen



For now, select Just For This Session. If you decide that you always want to use a certain window manager, you can click Make Default here the next time you use the Sessions menu, or change it permanently by creating a special configuration file in your home directory (which I'll explain a little later in this section). After clicking Just For This Session, the command files that start the X Window system on your computer will start the selected window manager. Figure 5-41 shows the IceWm window manager running on one of my Ubuntu systems.

FIGURE 5-41

The IceWM window manager on an Ubuntu system



To permanently change to using a specified window manager, you can either use the Sessions menu shown in Figure 5-39 and the subsequent dialog shown in Figure 5-40 to make a window manager your permanent choice, or use a text editor to create a file called .xsession in your home directory, which contains a line like the following:

```
exec /usr/bin/icewm
```

Make the file executable (chmod 755 ~/.xsession), and the next time that you log in on your system, the system will automatically start the IceWM window manager (in this example) instead of the GNOME desktop. For compatibility with other ways of starting the X Window system, you should also create a symbolic link to this file called .xinitrc, also in your home directory. You would do this with a command like the following:

```
ln -s ~/.xsession ~/.xinitrc
```

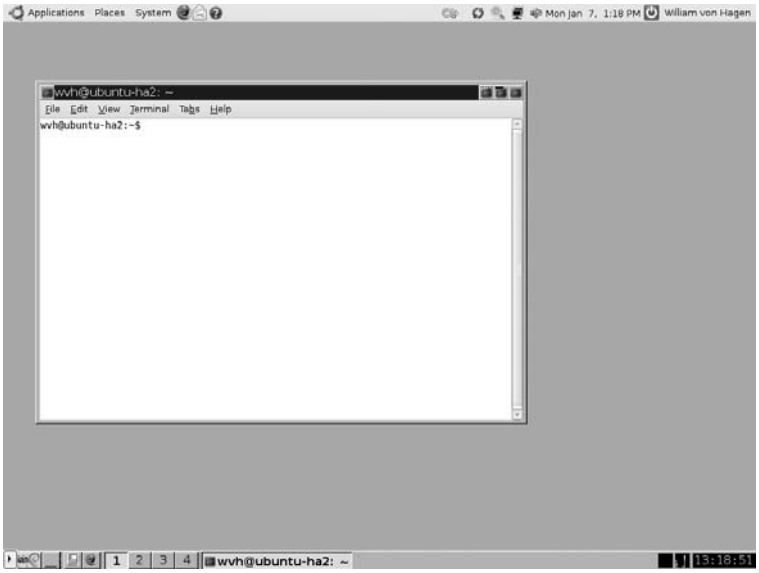
Experimenting with different window managers is fun because each represents someone's idea of how you "should" interact with a graphical system, start applications, and so on. I've been using the X Window system since version 10, and thus have burned more time than I could ever calculate in experimenting with, configuring, and customizing many of the window managers that have been written in the last 20+ years. But it's been fun.

TIP **On low memory, low resource systems such as older laptops, you may want to start Linux without a graphical interface — if you do this, you can always start the X Window system using a command such as `startx` or `xinit`, which will also pick up the X11 startup file and symbolic link that you just created, and will therefore start your window manager for you. See Chapter 22, "Understanding the Ubuntu Startup and Shutdown Processes," for information about changing whether your system starts in graphical or command-line mode.**

The previous instructions changed whether your Ubuntu system runs GNOME or a window manager. However, as mentioned earlier, all Linux desktop systems run a window manager under the covers. By default, GNOME on Ubuntu runs Compiz, KDE runs kwm, and so on. If you still want to run GNOME but are interested in changing the window manager that it runs underneath, you must run another window manager that conforms to the Inter-Client Communication Conventions Manual (ICCCM) 2.0 standard (`http://en.wikipedia.org/wiki/ICCCM`) published by the X Consortium, and also to the EWMH (Extended Window Manager Hints) guidelines. The former is an official specification for how X Window system clients running on the same server should interact, while the latter is an unofficial set of extensions to ICCCM 2.0 that were defined by the FreeDesktop.org folks to specify additional window manager, window, and utility interaction mechanisms. Window managers that conform to these specifications are often simply referred to as GNOME-compliant. Some GNOME-compliant window managers are After Step, Compiz, Enlightenment, Fvwm2, IceWM, Metacity, and WindowMaker. Figure 5-42 shows GNOME using the IceWM window manager.

FIGURE 5-42

GNOME running with the IceWM window manager



GNOME uses a variety of ways to change the window manager, some more complex than others. The simplest, which is something of a hack, is to modify your ~/.xsession file so that it first starts the window manager that you want to use, and then starts GNOME. The following is a sample ~/.xsession file that starts the IceWM window manager, and then starts GNOME:

```
/usr/bin/icewm&
exec gnome-session
```

This is a hack because it relies on the fact that only one window manager can be active at a time. By starting IceWM before starting GNOME, the Compiz window manager cannot start, and thus GNOME continues to run using IceWM.

You should actually be able to change your window manager by simply putting a line like the following in your ~/.gnomerc file:

```
export WINDOW_MANAGER=/usr/bin/icewm
```

Unfortunately, this doesn't work for me. If you search the Web, you'll find many other suggestions, such as modifying your gconf (GNOME configuration) database (using database entries that are explicitly unsupported in the latest GNOME releases), modifying GNOME window manager startup scripts (such as gnome-wm), modifying your session control file to reference the GNOME window manager scripts, and many combinations of these. In general, I'm a fan of things that "just work," so I prefer my "solution."

Summary

This chapter introduced the graphical user interfaces used on Linux systems, focusing on the GNOME desktop that is the default user interface provided on desktop Ubuntu Linux systems. After discussing the general organization of the GNOME desktop, the majority of the chapter focused on explaining how to customize and extend the default configuration of the GNOME desktop as installed on your Ubuntu system. Personalizing your computer system's desktop is one of the first steps to making it truly yours. GNOME provides the same types of customization and supports the same level of personalization as provided by other operating systems, such as Microsoft Windows or Mac OS X — it just doesn't cost as much. If you are using GNOME and are interested in customizing the Compiz window manager that is used by GNOME on Ubuntu desktop systems, see Chapter 6, "Using the Compiz Window Manager." If you are using Kubuntu or want to use the KDE desktop on a standard desktop Ubuntu system, Chapter 7 provides basic information about using and customizing version 3.*x* of KDE, while Chapter 8 discusses the latest (and possibly greatest) version of KDE, KDE 4.

If you're interested in learning about underlying Linux commands and using the Linux command line, Chapter 14 complements the exploration of the graphical power of GNOME, KDE, Compiz, and the X Window system in general by exploring the power of the Linux command line. The fact that many Linux commands are command line–oriented is one of the reasons that many hard-core Microsoft Windows users disparage Linux, which is just plain silly, because Linux gives you both a powerful graphical interface *and* a powerful command-line interface. Newer versions of Windows, such as Vista, have even added a "power shell" to enable Windows users to more easily perform various tasks from the command line. You don't have to look any further than Mac OS X, with its UNIX underpinnings and Terminal window, to see just how empowering the marriage of a great GUI and a powerful set of command-line tools can be in a commercial product. Ubuntu gives you this same power, if not more, but at no cost.

Chapter 6

Using the Compiz Window Manager

Among the most exciting recent developments in Linux graphics has been the advent of better-performing, open source graphics subsystems for the accelerated 2D and 3D graphics now provided by most modern graphics cards. Traditionally, only the specific X Window system graphics drivers that were delivered with more powerful video cards enabled users to take advantage of the high-end capabilities of these cards, and many graphics vendors didn't provide Linux drivers. Even when they did, the drivers themselves were typically precompiled and proprietary, and their source code was not freely available.

Today, however, life is better for users of more powerful graphics hardware under Linux. The use and popularization of standards such as OpenGL make it possible for open source systems to benefit from the capabilities of modern graphics cards, even with proprietary drivers. OpenGL (Open Graphics Library) is a standard that defines a cross-language, cross-platform API for writing graphical applications, providing hundreds of function calls in a standard API that can be used to draw complex 2D and 3D graphics. With the introduction of drivers that adhere to this standard, even proprietary video drivers that are available only in a binary format can be accessed and used from open source graphics applications and servers. Even better, the trend toward open source graphics drivers is accelerating, with open source drivers either being provided directly from graphics hardware vendors or developed by the open source community.

This chapter discusses the Compiz compositing window manager that has been installed by default on Ubuntu systems beginning with Ubuntu 7.10 and is also available for Kubuntu. After providing background information on exactly what compositing means and how compositing window managers work, this chapter explains how to configure and customize Compiz Fusion to provide sexy, high-performance graphics on your Ubuntu and Kubuntu systems.

IN THIS CHAPTER

- What is compositing?
- Compiz history
- Enabling special effects
- Dock software for Compiz
- Using Compiz with KDE

What Is Compositing? Why All the Fuss?

Most current graphics cards provide at least 128MB of onboard video RAM, while most of the video controllers integrated into today's motherboards enable you to allocate at least a similar amount of system memory to the graphics controller. Having a significant amount of memory available enables the video cards themselves to do substantial amounts of onboard graphics processing. Not only do today's video cards provide high-performance, built-in graphics processors, but offloading graphics calculations from the CPU and window manager to the video card improves overall system performance and responsiveness.

As discussed in Chapter 5, Linux systems use the X Window system (also known as X11, or simply X) to provide the graphical framework that underlies most Linux desktop environments and window managers. A window manager handles creating, moving, and managing windows shown on bitmapped displays that are managed by an X Window server, and they can be used either standalone or as components of a desktop environment that provides higher-level capabilities such as drag and drop.

The term "compositing" refers to the combination of graphical or visual elements from separate sources into single images. On Linux systems, a compositing manager is an X Window system client that redirects windows and other parts of the bitmapped display to video memory, tracks changes to that off-screen memory, and handles merging the updated images into the system's bitmapped display when changes occur and the screen has to be redrawn. An X Window system compositing window manager is therefore a unified window and compositing system that performs conventional window management functions while integrating and displaying high-quality, high-performance graphics being created and manipulated in your system's video memory, including special effects such as transparency and translucency.

TIP **Although the terms "transparency" and "translucency" are often used interchangeably, they actually have subtly different meanings for truly hardcore graphics fans. A transparent physical material shows objects that are behind it and doesn't reflect light off its surface. A translucent physical material shows objects that are behind it, but those objects are partially obscured because the translucent material reflects some of the light that hits it. See the OpenGL FAQ (`www.opengl.org/resources/faq/technical/transparency.htm`) for a more verbose explanation.**

Modern compositing window managers such as Compiz use separate applications to handle window decorations (window borders, buttons, frames, and so on), while the window manager focuses on managing window placement, movement, and special effects. This provides a flexible environment that follows the traditional UNIX philosophy of writing smaller applications that do one thing, do it well, and can work together to perform complex tasks.

Compiz, which is all the rage in Linux circles nowadays, provides a nice blend of making the best possible use of high-performance graphics hardware, providing attractive window effects for fans of eye candy, and providing a flexible window management environment that is easily modified and extended.

Compiz History: Compiz, Emerald, Beryl, and Friends

This section discusses the history of compositing window managers on Linux and the window managers that led to Compiz and Compiz Fusion today. Not everyone wants a history lesson, so if you're not interested, feel free to skip ahead. However, this section will provide a handy reference after you've heard terms such as Beryl, Emerald, Compiz, and Compiz Fusion over and over and wonder what they mean.

Like almost all graphics components on Linux systems, compositing window managers depend on the capabilities of the underlying X Window system. To provide high-performance special effects, compositing window managers require especially tight integration between graphics drivers and the window manager. Compiz is not the first compositing window manager — compositing has been supported for a while in window managers such as Metacity, the window manager that was traditionally associated with GNOME 2.2 (and used on Ubuntu releases prior to 7.10), and Xfvwm, the window manager used by the Xfce desktop (`http://xfce.org`). As will be discussed in Chapter 7, compositing support is also available in KWin, the window manager used by the KDE desktop and therefore Kubuntu systems, beginning with KDE version 4.0.

Regardless of the fact that Metacity, Xfvwm, and some lesser-known window managers pioneered built-in support for compositing, Compiz was the first widely available window manager that was designed from the ground up to support compositing via OpenGL. The efficient support of OpenGL requests between an X Window server is handled using one of two technologies, depending upon the X Window system server in use:

- X Window servers such as Xgl require extremely tight communication between the X Window server and the graphics hardware, providing impressive performance on a more limited range of graphics cards from vendors such ATI, Intel, and NVIDIA.
- X servers such as the standard X.org server (as of release 7.1) use a more device-independent mechanism known as AIGLX (Accelerated Indirect GLX) that enables them to work with a wider range of graphics hardware.

Ubuntu systems use the standard X.org window manager and thus use the AIGLX approach.

TIP **The only reason I'm discussing Xgl/AIGLX here is that I've found that some systems with ATI or NVIDIA graphics chips work better with the Xgl server than with the standard X.org server. This is especially true for laptop computers, which often use ATI or NVIDIA chips that require proprietary or highly specialized drivers. If you have problems getting Compiz working with the standard X server on a system with ATI or NVIDIA chips, it's worth trying the Xgl server.**

Compiz was designed to provide a flexible window management framework to which plug-ins could easily be added to provide additional capabilities and effects without the need to reinstall the entire window manager. Compiz and Xgl were both originally developed at Novell, and thus Compiz originally required Xgl. Thanks to AIGLX and lots of hard work, Compiz now also works with the standard X.org server, and the AIGLX development effort has been merged directly into X.org.

One of the reasons behind the foundation of the AIGLX effort was that some people felt that too much of the development of Xgl was local to Novell and didn't invite enough community discussion. Similarly, the Beryl project emerged as a fork of Compiz when people felt that the development of Compiz was moving too slowly. Beryl was previously known as the Quinnstorm version of Compiz, but as the codebases diverged, it seemed more reasonable to maintain two separate projects.

As mentioned in the previous section, one of the core ideas of Compiz is to use separate applications to handle window decorations. Compiz provides the `gtk-window-decorator` to customize window decorations in the GNOME environment (and can even use standard Metacity themes), and the `kde-window-decorator` uses native KDE themes. Beryl's primary window decorator is Emerald, although window decorators named Heliodor and Aquamarine were available at one time.

Beryl racked up a tremendous set of plug-ins and whizzy features, while Compiz moved to AIGLX and became more open. Duplication of effort is rarely a good thing, and in 2007, the people working on the Beryl project merged with the folks who were developing Compiz plug-ins, resulting in the Compiz

Fusion project. The goal of this project is to improve the usability and functionality of Compiz by porting most of the whizzy things that were done for Beryl to Compiz, and to continue to develop additional libraries, plug-ins, and tools that make Compiz both a fun and functional graphical environment. The Compiz Fusion folks have even ported Emerald to work with Compiz, to ensure that you will always have an impressive (and sometimes bewildering) smorgasbord of graphical choices. Today's Compiz comes with a core set of plug-ins, but the plug-ins and tools provided by Compiz Fusion make the whole far greater than the sum of its parts.

TIP **If you are having problems running Compiz or its performance seems sluggish, consult the Compiz Fusion Web site (`http://wiki.compiz-fusion.org/Troubleshooting`) for good information on identifying and resolving such problems.**

Enabling Special Effects in Compiz

As mentioned in the previous section, Compiz provides a flexible and easily extended compositing window management environment. Much of the functionality that people lump together as Compiz is actually provided by plug-ins that work within the framework provided by Compiz. Similarly, Compiz enables you to run separate window decorator applications to provide more granular control over such decorations, as well as over window controls and themes.

NOTE **When switching to Compiz, you may find some of the terminology confusing, particularly *workspace*, *virtual desktops*, and *viewports*. A workspace is a shared area that one or more applications can access — in the case of a window manager, think of it as a distinct graphical region on which windows can be displayed. Workspaces are independent of one another and have their own graphical coordinate systems. A viewport is a view of a graphical region on the screen. Traditionally, GNOME and its window managers used multiple workspaces to support virtual desktops; each virtual desktop was its own workspace. In Compiz, each virtual desktop is a viewport into a shared workspace, and the use of a viewport improves performance and simplifies various graphical operations. For example, when you are using workspaces, each window should exist within only a single workspace. When you are using viewports, windows can be visible in multiple viewports, which provides for smoother graphic transitions from one virtual desktop to another and also simplifies implementing edge resistance, fast virtual desktop rotation and switching, and so on.**

Using the Appearance Preferences Application

By default, when Compiz is installed, a basic set of plug-ins is activated. The most basic configuration utility for controlling the special effects associated with Compiz is the Appearance Preferences application, which you can start by selecting Appearance from the Preferences menu. After starting this application, select the Visual Effects tab, which displays the configuration options shown in Figure 6-1.

The following are the three basic appearance preferences settings:

- **None:** Selecting this option disables any special effects associated with Compiz and Compiz Fusion plug-ins, providing a Compiz desktop experience that is as near as possible to the Metacity desktop experience provided on Ubuntu releases prior to version 7.10.

TIP **As noted, selecting None will disable Compiz and Compiz Fusion plug-ins. You will not be able to enable or use these plug-ins, even if you manually enable them using one of the Compiz configuration utilities discussed in the next section. If you are having difficulty getting any Compiz or Compiz Fusion plug-ins to work correctly, first check that you have not accidentally disabled plug-ins by selecting this Appearance Preferences option.**

- **Normal:** Selecting this option, which is the default setting if compositing could be enabled during installation, activates a small but useful set of Compiz and Compiz Fusion plug-ins. If your system could not enable compositing during installation (which depends on your video hardware), selecting this option may require that your system search for new video drivers. Once this option is active, a number of Compiz visual effects are automatically enabled. In my opinion, the most interesting of these are the following:
 - **Animations:** Provides graphical effects when performing window operations such as minimizing and maximizing applications, starting and terminating applications, and so on.
 - **Application Switcher:** Provides a more attractive Alt+Tab interface (for switching between open applications), advanced image loading, and text rendering effects.
 - **Desktop Wall:** Provides a smooth virtual desktop/viewport–switching application.
 - **Expo:** Displays reduced-size images of all of your windows on the screen so that you can quickly select a specific window.
 - **Move Window:** Provides mouse and keyboard shortcuts for moving windows, which can be handy if the title bar is not easily selectable.
 - **Resize Info:** Provides a shaded preview when resizing windows and displays numeric information about the current size of the new window as you move the mouse.
 - **Snapping Windows:** Enables you to move a window to other desktops by dragging a window to the edge of the screen, and lets you configure how quickly this happens to prevent you from moving windows to other desktops accidentally.
 - **Window Decorations:** Additional window decoration features such as shadowing (required for some window decorators, such as Emerald).

FIGURE 6-1

Basic effects management via Appearance Preferences



The Normal setting also enables several other plug-ins, most of which perform internal functions or are not bound to specific keys by default.

- **Extra:** Selecting this option activates the plug-ins associated with the Normal Appearance Preferences settings and activates other plug-ins, the most interesting of which are the following:
 - **Shift Switcher:** An enhanced version of the Application Switcher, this plug-in displays available windows as minimized, angled previews that you can switch between using the Super+Tab key sequence rather than the traditional Alt+Tab sequence.
 - **Wobbly Windows:** One of the best-known Compiz effects, this plug-in causes the window borders to jiggle as you move a window, making it look as if your window is made out of Jell-O as it's being moved (fun!)

TIP

Most Compiz plug-ins are associated with specific actions, key combinations (known as keyboard shortcuts), and/or mouse movements. While these are configurable, most of the default key combinations require the use of the Super key on your keyboard, which is usually the key with the Windows logo that is positioned between the Ctrl and Alt keys on most newer keyboards. If your current keyboard doesn't provide this key, I suggest getting a new keyboard that does (unless, of course, you are using an older laptop). You can always remap the keys used to perform various actions, but investing in a new keyboard is generally preferable to investing the amount of time it would take you to change everything.

As you can see, the default sets of Compiz and Compiz Fusion plug-ins provided by the Normal and Extra Appearance Preferences sets provide a good deal of functionality in a graphically attractive fashion. However, Compiz and Compiz Fusion provide many other plug-ins, some of which are alternatives to the ones that are already provided by the Normal and Extra configurations, and others that provide additional capabilities. The list of Compiz and Compiz Fusion plug-ins is always expanding, so it is impossible to give you a complete list here. The sidebar "These Are a Few of My Favorite Things," later in this chapter, lists some other plug-ins that I find useful and that I think you might also find useful and enjoyable. For a complete listing of available Compiz and Compiz Fusion plug-ins and information on their configuration, see `http://wiki.compiz-fusion.org/Plugins`.

Detailed Compiz Configuration

The configuration mechanisms discussed in the next two sections enable you to activate, customize, and deactivate Compiz and Compiz Fusion plug-ins, giving you access to the complete list of plug-ins that are currently installed on your system and all of their configuration settings.

On Ubuntu systems, you can configure Compiz, the Compiz plug-ins, and the Compiz Fusion plug-ins by using the standard GNOME configuration utilities or by using the CompizConfig Settings Manager. Other Compiz configuration utilities are also available, but GConf and CompizConfig are the two biggies. Of these two configuration mechanisms, I strongly recommend installing and using the CompizConfig Settings Manager, but the standard GNOME utilities will work just as well if you're a fan of the command line or already intimately familiar with the graphical GNOME Configuration Editor. However, I'll discuss both approaches to help you make your decision.

Configuring Compiz Using GNOME Configuration Tools

On Ubuntu systems, Compiz, the Compiz plug-ins, and the Compiz Fusion plug-ins can be configured using the standard GNOME configuration utilities without requiring that you install any special tools.

GNOME systems provide two primary configuration utilities for manipulating the GConf (GNOME Configuration Editor) repository: `gconftool-2`, a command-line utility, and `gconf-editor`, a graphical application. You should run these utilities as yourself, without using `sudo`, because they manipulate GNOME configuration information stored in each user's personal GConf repository. If you ran these utilities via `sudo`, you'd edit the GConf repository for the root user.

Information in the GConf repository is stored as key/value pairs that are organized into various hierarchical sections, much like the information in the Microsoft Windows Registry. Information about Compiz is stored in the `/apps/compiz` hierarchy. Information about specific plug-ins is hierarchically organized under `/apps/compiz/plugins` and the name of each plug-in.

`gconftool-2` is a command-line line application that enables you to examine all or a portion of the GConf registry and set specific values by passing repository keys and new values on the command line. For example, the following command examines the settings for the Compiz Wobbly Windows plug-in:

```
$ gconftool-2 -R /apps/compiz/plugins/wobbly
  /apps/compiz/plugins/wobbly:
   /apps/compiz/plugins/wobbly/allscreens:
    /apps/compiz/plugins/wobbly/allscreens/options:
     shiver = false
     snap_key = <Shift>
     snap_inverted = false
   /apps/compiz/plugins/wobbly/screen0:
    /apps/compiz/plugins/wobbly/screen0/options:
     spring_k = 8
     map_effect = 0
     map_window_match = Splash | DropdownMenu | PopupMenu | \
          Tooltip | Notification | Combo | Dnd | Unknown
     focus_window_match =
     friction = 3
     focus_effect = 0
     maximize_effect = true
     grab_window_match =
     min_grid_size = 8
     grid_resolution = 8
     move_window_match = Toolbar | Menu | Utility | Dialog | \
          Normal | Unknown
```

The `-R` option tells `gconftool-2` to recursively retrieve hierarchical settings below the specified repository key. Information about whether a given plug-in is enabled is stored in the `/apps/compiz/general/active_plugins` repository entry, as shown in this example:

```
$ gconftool-2 -g /apps/compiz/general/allscreens/options/active_plugins
[svg,imgjpeg,text,place,png,neg,workarounds,dbus,video,regex,\
resize,core,vpswitch,move,resizeinfo,shift,extrawm,decoration,\
wall,animation,expo,ezoom,switcher,fade,scale,scaleaddon,scalefilter]
```

NOTE **As with all examples of command-line utilities in this book, the \ character is used at the end of text lines that are too long to be correctly displayed in the running text of this book. You would not see this character in actual output, and you should not use this character, the subsequent newline, or any spaces at the beginning of the next line when actually typing commands, as shown in the examples.**

The `-g` option retrieves the value of the specified repository key. The output from this command shows that the Wobbly Windows plug-in is not currently enabled.

To activate the Wobbly Windows plug-in from the command line, you would use the `-s` option, which lets you set the specified repository key to a specified value. The `-s` option must be used with the `--type` option to identify the type of the value that you are setting. Valid options for the `--type` option are `bool` (Boolean), `float` (floating point numeric value), `int` (integer), `list` (a list of values of types `bool`, `float`, `int`, or `string`), `pair` (two distinct values), or `string` (a string of text). Identifying the type of value being set is required so that `gconftool-2` can determine, for example, whether the value `true` is the string `true` or a Boolean value.

The command to set the `active_plugins` repository key to a value that activates the Wobbly Windows plug-in is something like the following:

```
$ gconftool-2 -s /apps/compiz/general/allscreens/options/active_plugins \
'[svg,imgjpeg,text,place,png,neg,workarounds,dbus,video,regex,resize,\
  core,vpswitch,move,resizeinfo,shift,extrawm,decoration,wall,animation,\
  expo,ezoom,switcher,fade,scale,scaleaddon,scalefilter,wobbly]' \
  --type list --list-type string
```

You must begin and end lists with the characters `[` and `]` in order to delimit the lists. Once you execute this command, reexamining the value of the `active_plugins` key shows that the Wobbly Windows plug-in is now in the list of active plug-ins:

```
$ gconftool-2 -g /apps/compiz/general/allscreens/options/active_plugins
[svg,imgjpeg,text,place,png,neg,workarounds,dbus,video,regex,resize,\
 core,vpswitch,move,resizeinfo,shift,extrawm,decoration,wall,animation,\
 expo,ezoom,switcher,fade,scale,scaleaddon,scalefilter,wobbly]
```

You can also confirm that this plug-in is enabled by using your mouse to grab the title bar of any window and move that window. There's always room for Jell-O!

`gconftool-2` can be very useful for quickly querying and manipulating GConf settings from the command line. However, `gconftool-2` is probably not the right tool for setting the "right" GConf keys, identifying possible values, and iterating through multiple settings and values until you get the effect that you're looking for. I'm a big fan of command-line tools, but I'm no purist — in most cases, graphical tools are more usable than command-line tools, especially when you need to perform a number of change and test cycles when experimenting with configuration values.

GNOME's graphical tool for examining and editing the GConf repository, `gconf-editor`, is installed by default on GNOME-based Ubuntu systems. You can start the `gconf-editor` tool from any command-line environment. As with `gconftool-2`, you want to start `gconf-editor` as yourself, not via sudo, so that you are editing only the contents of your personal GNOME repository.

Figure 6-2 shows `gconf-editor` when it's first started. To locate a specific repository key, click the Edit ➪ Find command to display the dialog shown in Figure 6-3.

Enter the name of the key that you want to search for (`active_plugins` in the figure) and select the Search also in key names option to expand the types and places that are being searched for the specified key. By default, `gconf-editor` looks for only the search string you specify in key values. In this case, you want to search key names because the list of active plug-ins is stored as the value of the `/apps/compiz/general/active_plugins`. After you click Find, you'll see a screen like the one in Figure 6-4,

in which the matching keys are listed in the bottom pane. Select the `/apps/compiz/general/active_plugins` key to display the view shown in Figure 6-4.

FIGURE 6-2

The graphical GNOME configuration editor

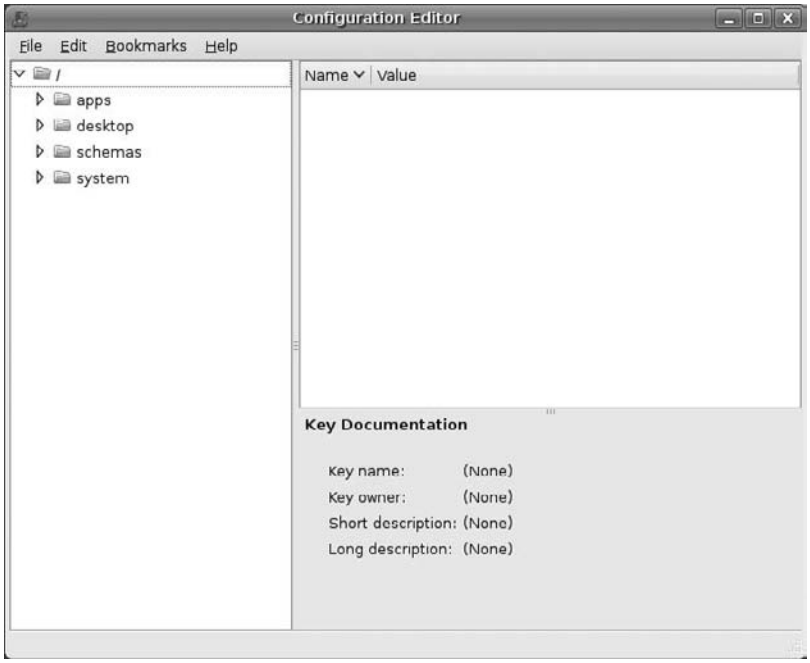


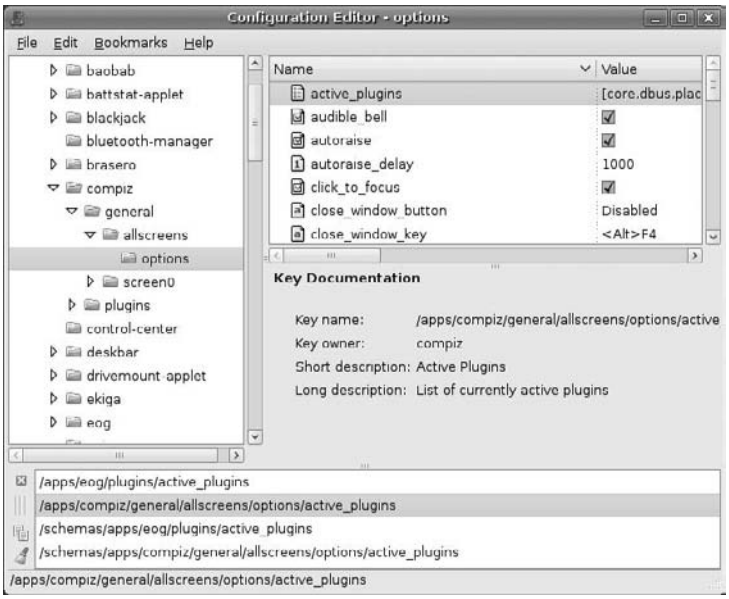
FIGURE 6-3

Searching for a repository key



FIGURE 6-4

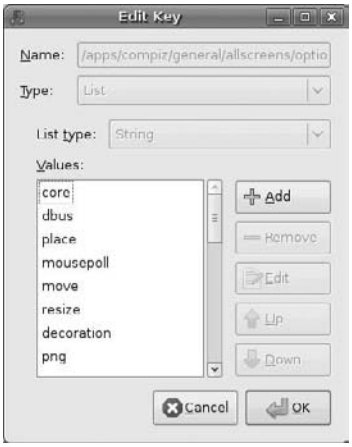
Repository matches for a search



To edit an existing key, right-click the name of that key and select Edit key from the pop-up menu. You'll see the dialog shown in Figure 6-5.

FIGURE 6-5

Editing a key value



As you can see from Figure 6-5 (and from the examples earlier in this section), the value of the `/apps/compiz/general/active_plugins` key is a list of plug-in names. If the Wobbly Windows plug-in is not present in this list (which it is not by default, but you may have added it earlier in this section if you're following along), click the Add button to display the Add New List Entry dialog shown in Figure 6-6.

FIGURE 6-6

Adding a new entry to a key value list



Enter **wobbly** as the new list value and click OK. The Wobbly Windows plug-in will be active at this point. You can then select File ➪ Quit to close the `gconf-editor` application.

Both the `gconftool-2` and `gconf-editor` tools provide a generic interface to the GNOME configuration registry. While these tools provide a completely functional mechanism for examining, adding, and modifying repository key values, they simply see the repository as a hierarchical set of key/value pairs and thus see every configuration task as being essentially equivalent.

Unlike these standard GNOME registry configuration tools, the tool discussed in the next section was designed for manipulating Compiz and Compiz Fusion plug-ins, and can therefore make use of plug-in–specific information to provide an easier-to-use, smarter plug-in configuration tool.

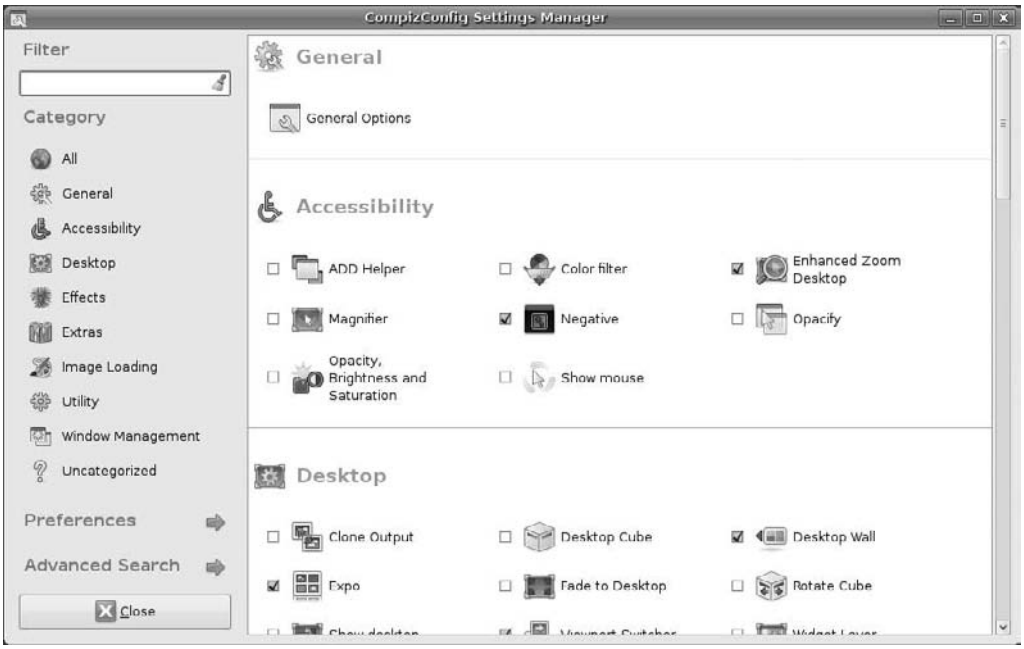
Installing and Using CompizConfig Settings Manager

The CompizConfig Settings Manager, sometimes referred to as CCSM by fans of abbreviations, is a Compiz configuration utility created by the Compiz Fusion project (`http://wiki.compiz-fusion.org/CCSM`). Why this configuration utility isn't installed by default with Compiz on Ubuntu systems is a mystery to me, because it provides an excellent interface that simplifies activating, configuring, and deactivating both Compiz and Compiz Fusion plug-ins. CompizConfig was designed solely for the purpose of configuring Compiz plug-ins. It is able to make use of some information about the internals of Compiz plug-ins and the capabilities that they provide and can therefore identify both dependencies and potential conflicts between plug-ins.

You can install CompizConfig by using your favorite package management tool to install the `compizconfig-settings-manager` package, as described in Chapter 23. This package is located in the Universe repository, so that repository must be enabled in your `/etc/apt/sources.list` file before you can install the package. (By default, the Universe repository is enabled in this file, but you may have disabled it for some reason.) Once you've installed the CompizConfig package, you can start the CompizConfig Settings Manager by selecting Preferences ➪ Advanced Desktop Effects Settings. Figure 6-7 shows the CompizConfig Settings Manager when you first start it.

FIGURE 6-7

The CompizConfig Settings Manager



Most configuration utilities make it incredibly easy for you to iteratively change too many options, to the point where you don't know how to restore things to their original values (or at least know what those original values were). The CompizConfig interface provides a nice, consistent solution to this problem by providing a small icon that looks like a whisk-broom to the right of each individual configuration setting. Clicking this icon at any time restores the associated configuration option to its original, default value.

To enable any plug-in but use its default configuration, click the checkbox to the left of the icon and name for that plug-in, as shown in Figure 6-7. To configure any Compiz or Compiz Fusion plug-in, click the name or icon of that plug-in. This displays the configuration panel for that plug-in, an example of which is shown in Figure 6-8. You can also enable or disable the plug-in after configuring it by selecting or deselecting the checkbox below the Use This Plugin label in the left-hand pane of a configuration panel such as the one shown in Figure 6-8. You can return to the top-level CompizConfig panel by clicking the Back button shown in the lower-left corner.

The General Options item at the top of the CompizConfig screen (refer to Figure 6-7) enables you to configure options for Compiz itself, and doing so may affect multiple (or all) plug-ins. The General Options configuration item is organized into seven tabs, as shown in Figure 6-9.

FIGURE 6-8

Configuring a specific plug-in in CompizConfig

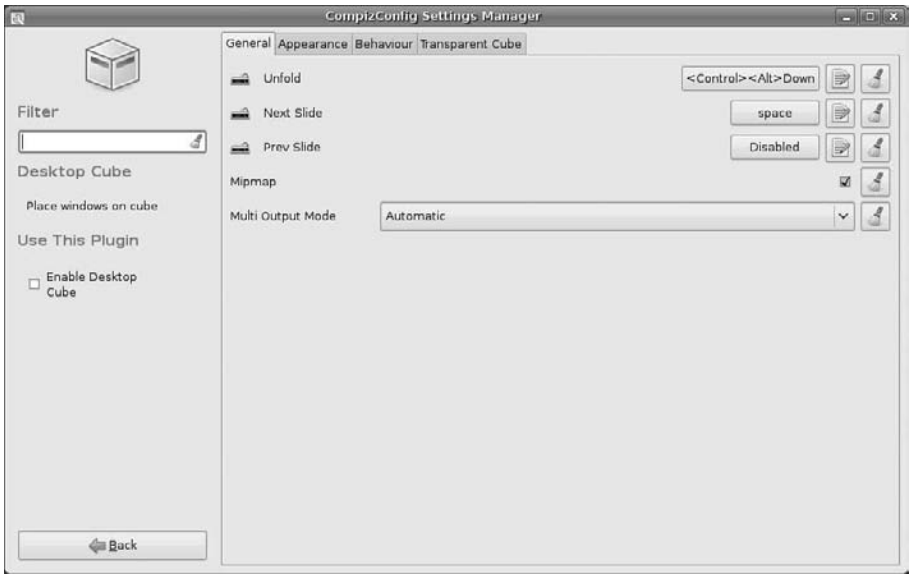
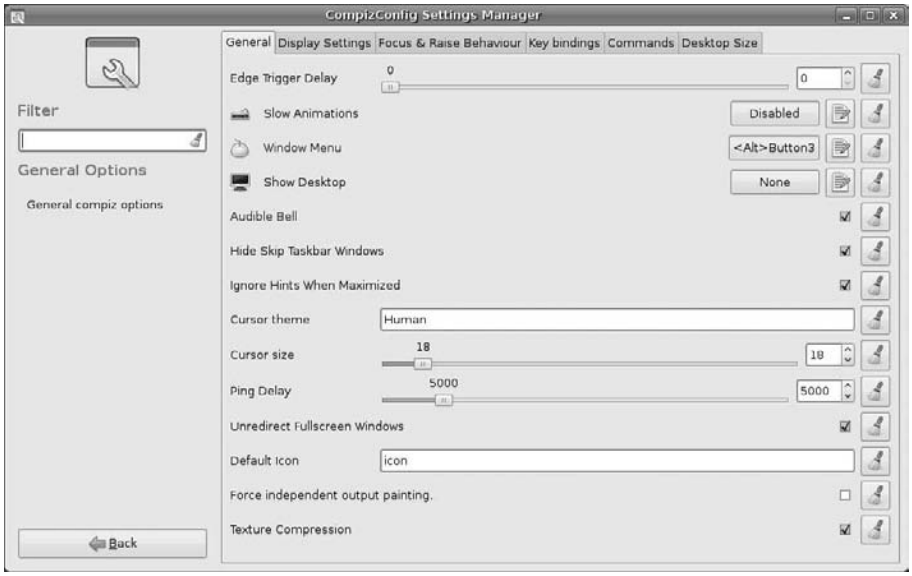


FIGURE 6-9

Setting General Options in CompizConfig



Each of these tabs contains multiple configuration options, such as:

- **General:** Enables you to configure the cursor and its size, configure the terminal application for use with command-line functions, control whether Compiz provides audible notification of warning and error messages, and so on.
- **Display Settings:** Enables you to configure general aspects of how Compiz interacts with your display: how finely the screen displays surface textures, whether Compiz attempts to automatically identify your screen's refresh rate, whether the screen or a window dims as it is transformed (Lighting), the screen's refresh rate, whether screen updates are limited to taking place during vertical blanking, the output devices that Compiz knows about and whether it should attempt to identify them automatically, and so on.
- **Focus & Raise Behaviour:** Enables you to specify how your mouse interacts with and selects windows: whether you have to click a window to set the focus to it, whether windows rise to the top automatically when you position the mouse over them (and how long it takes for this to happen), whether clicking anywhere in a window brings that window to the top, and which windows will not automatically steal the focus when they appear. (The last of these requires that the Regex Matching plug-in be enabled, as does any plug-in that needs to be able to identify specific windows or types of windows.)
- **Key Bindings:** Enables you to define the key bindings, mouse actions, or screen regions associated with common window operations (General section) as well as the commands to run when specific `run_command` functions are invoked (the Commands section). The commands on this tab also enable you to map keys, mouse actions, or screen regions to the Increase Opacity and Decrease Opacity operations.
- **Commands:** Enables you to set the Linux commands that are executed when the internal Compiz `run_command` commands and the built-in screenshot commands are invoked. The former provides a generic way for you to map specific commands to specific keys. You can define the key mapping that executes each of the `run_command` commands in the Action tab's Commands section. For example, you could map the Super-1 key combination to automatically execute the `top` command within an `xterm` window by entering `xterm -e top` as `run_command1`, and then using the Action tab's Commands section to map `run_command1` to the key sequence Super-1.
- **Desktop Size:** Enables you to set the number of viewports that are supported by Compiz for each desktop, and their horizontal and vertical organization. As I'll discuss later in this section, the horizontal virtual size must be set to 4 if you are to use the Desktop Cube plug-in. Its default value is 2, which is confusing.

As an example of using CompizConfig, let's enable the Rotate Cube plug-in, which is one of the best-known Compiz plug-ins. When completely configured (along with some related plug-ins), this plug-in displays your active viewports as the four sides of a cube that you can quickly rotate to cycle through the available viewports. You can also rotate the cube in three dimensions to expose its top and bottom, which you can decorate with images. Why have an incredibly powerful graphics card if you're not going to make the most of it? To activate the Rotate Cube plug-in, follow these steps:

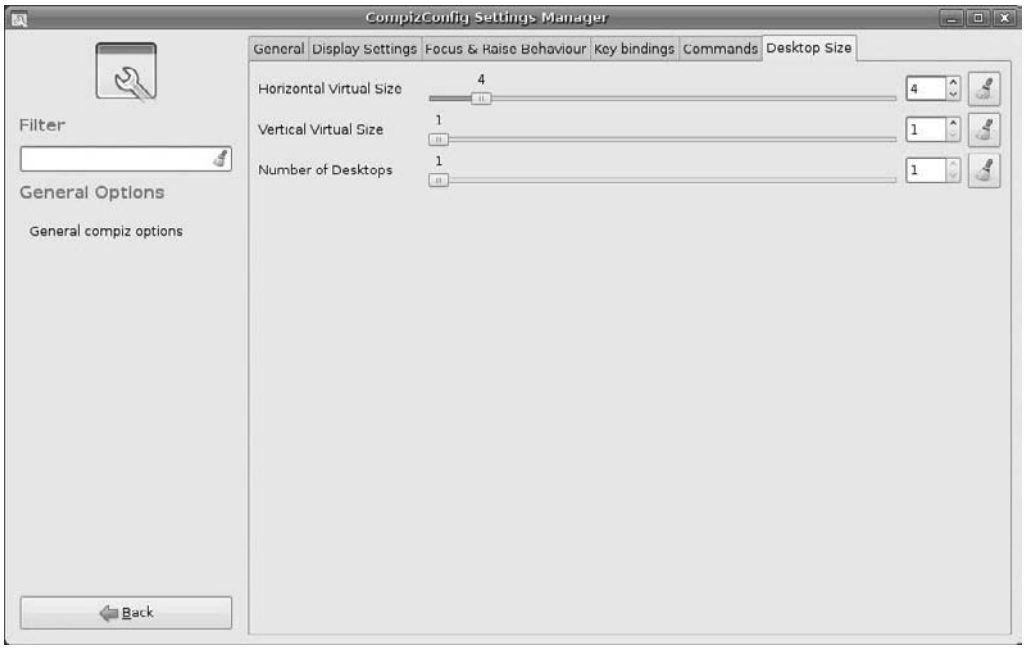
1. Configure your desktop so that four viewports are available.
2. As mentioned earlier in this section, you must use the General Options' Desktop Size tab to set the Horizontal Virtual Size of your desktop to 4 in order to provide 4 viewports that you can use as the four sides of your desktop cube, as shown in Figure 6-10. (The default value is 2, which is confusing.) To do this, select the General Options section of CompizConfig, select the Desktop

Size tab, and either enter the new value of 4 or use the up arrow to increase the default value of 2 for the Horizontal Virtual Size parameter.

When you're finished, your screen should look like that shown in Figure 6-10. Click the Back button to return to the main CompizConfig screen.

FIGURE 6-10

Increasing the horizontal viewports on your system



3. Activate the Desktop Cube plug-in by clicking the X to the left of its icon and name. If you have made no previous configuration changes, you will see a warning dialog like the one shown in Figure 6-11.

FIGURE 6-11

Identifying potential plug-in conflicts



Dialogs of this type are one of the advantages of using a domain-specific configuration applications such a CompizConfig instead of a generic configuration application such as `gconf-editor`. Because CompizConfig knows more about plug-ins than simply how to manipulate the GNOME registry, it can detect feature conflicts and offer to disable active plug-ins that might cause them to fail when you enable new plug-ins. In this case, click Disable Desktop Wall, which will disable the enhanced, but linear, viewport scroller and will enable the Desktop Cube plug-in, as requested.

4. Select the checkmark beside the Rotate Cube plug-in, which is the plug-in that actually rotates the cube (the Desktop Cube simply maps viewports to the sides of a cube but doesn't provide any whizzy graphics). At this point, you should be able to rotate the faces of your new desktop cube in one of three basic ways:

 - By using the Ctrl+Alt+Right key combination to rotate the desktop to the right, and the Ctrl+Alt+Left key combination to rotate the desktop to the left.
 - By clicking the title bar of a window and dragging that window to the left or right edge of the screen. The cube will rotate, and the window you're moving will be moved to the new viewport.
 - By positioning the mouse cursor on your desktop background and moving the mouse's scroll wheel (if you have one).

 Pretty impressive, no? Even more eye-catching are the following:

 - You can spin the cube in three dimensions by holding down the Ctrl and Alt keys, clicking and holding the primary mouse button, and moving the mouse up and down to expose the top and bottom of the cube.
 - You can unfold the cube (flattening it to a single screen) by pressing the Ctrl+Alt+Down keys, and you can then scroll left or right by using the appropriate arrow keys. This makes it easy to locate the viewport that contains the applications you're looking for. Once the viewport that you're looking for displays in the center of the screen, release the Ctrl, Alt, and arrow keys to display it full-size.

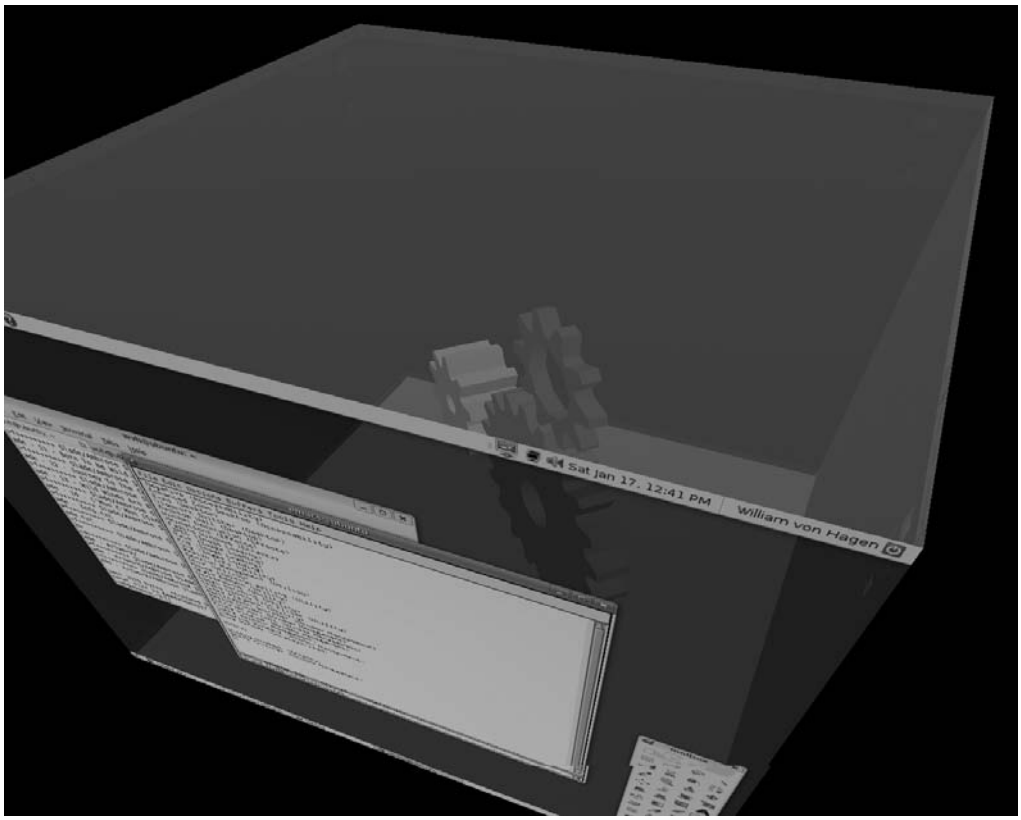
Now that you've had some fun with this plug-in, you may want to consider tweaking it even further. Some fun suggestions are:

- Use the Desktop Cube's Appearance tab to add custom images to the top and/or bottom of the desktop cube by specifying image files in the Cube Caps section.
- Use the Desktop Cube's Appearance tab to specify a background image or gradient for the space "behind" the 3D rotating cube by selecting a gradient start and end color or specifying image files in the Skydome section.
- Set the cube to be transparent using the Desktop Cube's Transparent Cube tab, and then activate the Cube Gears plug-in to display a rotating set of gears inside the cube as it rotates. Figure 6-12 shows a transparent cube running the cube gears plug-in in the process of being rotated.

These suggestions show the fun you can have with Compiz and its plug-ins and the incredible range and power of the graphical capabilities it provides. Given the number of potential settings you can configure and the bells and whistles that Compiz provides, you may never get around to doing real work again!

FIGURE 6-12

Rotating a transparent cube while using the cube gears plug-in



TIP **A common Compiz problem is that the title bars of newly-created windows display under the GNOME panel at the top of the screen. If this occurs, first make sure that the Place Windows plug-in is enabled. This plug-in uses a "smarter" algorithm that takes window title bars into account when placing new windows on the screen. You can also simply move such windows so that their title bars are visible by holding down the Alt key and clicking and holding the left mouse button in the window that you want to move. The standard window move icon displays, and you can then drag the window such that its title bar is visible.**

These Are a Few of My Favorite Things

Some of the plug-ins that I typically activate beyond the default set are as follows:

- **Desktop Cube/Rotate Cube:** Discussed in the previous section, quickly cycling through available viewports and rotating a cube with those viewports on its sides is not only attractive but also eminently useful.
- **Cube Gears:** Although I realize I'm using more of the CPU's resources, this is such a cool effect that I can't resist activating it.
- **Ring Switcher:** This plug-in adds the Super-Tab command to display available windows in a circular ring that spins as you tab through them, with the title of the selected window displayed in the center. It's similar to the standard Alt+Tab application Switcher but much more fun.
- **Water Effect:** Sure to amuse the kids and amaze the neighbors, this plug-in lets you toggle onscreen raindrops with Shift+F9. This wastes computing power only when the raindrops are active, but it's quite visually impressive when you're showing Compiz to someone.
- **Wobbly Windows:** Discussed in the section of this chapter on using the GNOME configuration tools to configure Compiz effects, Wobbly Windows provide a jiggly, gelatinous view as you move windows. Like the Desktop Cube/Rotate Cube plug-ins, this plug-in provides one of the classic Compiz effects.

As mentioned earlier in this chapter, see `http://wiki.compiz-fusion.org/Plugins` for a complete listing of available Compiz and Compiz Fusion plug-ins and their configuration. The Compiz plug-ins page at `www.compiz.org/Plugins` lists the default plug-ins that are bundled with Compiz and also provides links to some additional, external plug-ins that you may find useful.

Other Useful Packages for Compiz

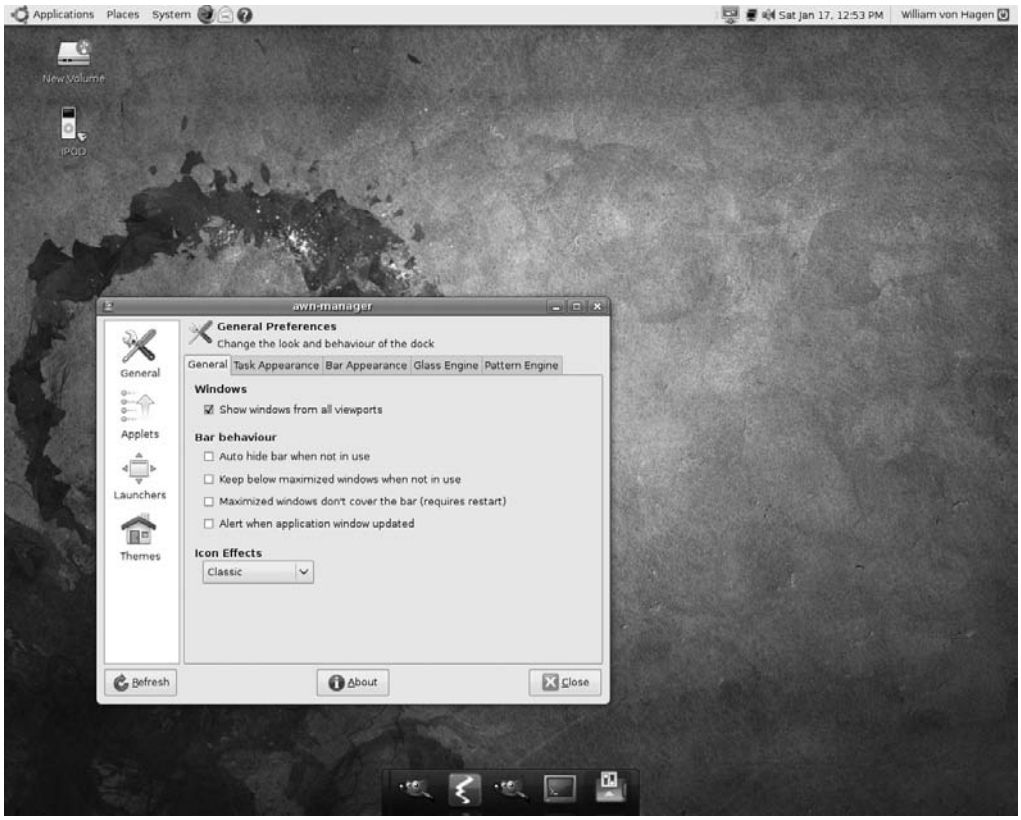
As the previous section showed, some extremely useful Compiz-related software is available in the Ubuntu repositories but is not installed by default on Ubuntu or Kubuntu systems. The previous section discussed the CompizConfig Settings Manager; this section and its subsections discuss two other very interesting packages that you can install for use with Compiz.

Installing and Using the Avant Window Navigator

If you've ever seen or used Mac OS X, you're probably familiar with its dock, which is the panel-like application that runs at the bottom of the screen (by default), and in which you see icons for your favorite applications as well as icons representing running applications and the Trash Can. Although GNOME has its panels, KDE 3 has the Kicker, KDE 4 has its Kickoff panel, and Microsoft Windows has its Taskbar, it's hard to deny that the Mac OS X Dock is the most attractive of any these entities with similar functionality. Perhaps until now. The Avant Window Navigator (AWN) is an open source dock application for compositing window managers that provides an attractive, powerful, extensible, and fully customizable dock for use with Compiz. The Avant Window Navigator was developed by Neil Patel, and currently has its home page at the Ubuntu Launchpad site at `https://launchpad.net/awn`. Figure 6-13 shows the default appearance of an Ubuntu desktop running both the Avant Window Navigator and its configuration utility, the Awn Manager.

FIGURE 6-13

The Avant Window Navigator and its administrative application



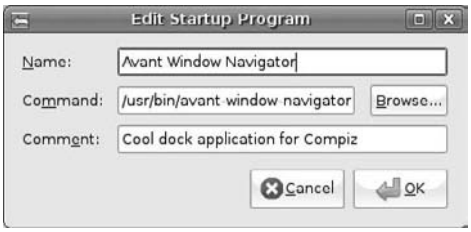
The Avant Window Navigator is not installed by default on Ubuntu systems, but you can easily install it using your favorite package management utility, as explained in Chapter 23. Installing the Avant Window Navigator will also install the library that it uses, some Python modules that it required, and its configuration manager.

NOTE **The Avant Window Navigator can also be installed on Kubuntu systems using the Adept tools, but because it is a GNOME-based application, installing it will also install a complete GNOME environment on your system. That's a lot of disk space for a dock, but it's completely up to you.**

After installing AWN, you can start it manually from any command-line application, such as an xterm, the GNOME Terminal, or the Konsole. To configure AWN to start automatically when you log in, start the System ➪ Preferences ➪ Sessions application. On the Startup Programs tab, click Add, and enter information about the Avant Window Navigator, as shown in Figure 6-14.

FIGURE 6-14

Configuring the Avant Window Navigator to start automatically at login



On Ubuntu systems, I prefer to eliminate the bottom panel and add the useful applets that it normally contains (most notably the WorkSpace Switcher) to the top panel, as shown in Figure 6-13. Using AWN at the bottom of the screen and the standard GNOME panel at the top gives me the best of both worlds.

Once you've installed and started AWN, you can start the Awn Manager application by selecting System ➪ Preferences ➪ Awn manager. This displays the application shown in Figure 6-13, which enables you to customize how icons in the dock are displayed, configure whether the dock is automatically hidden when not in use, configure the appearance of the bar underneath the dock, and add applets, launchers, and themes to customize the dock's behavior and appearance. Figure 6-15 shows AWN with the following customizations enabled:

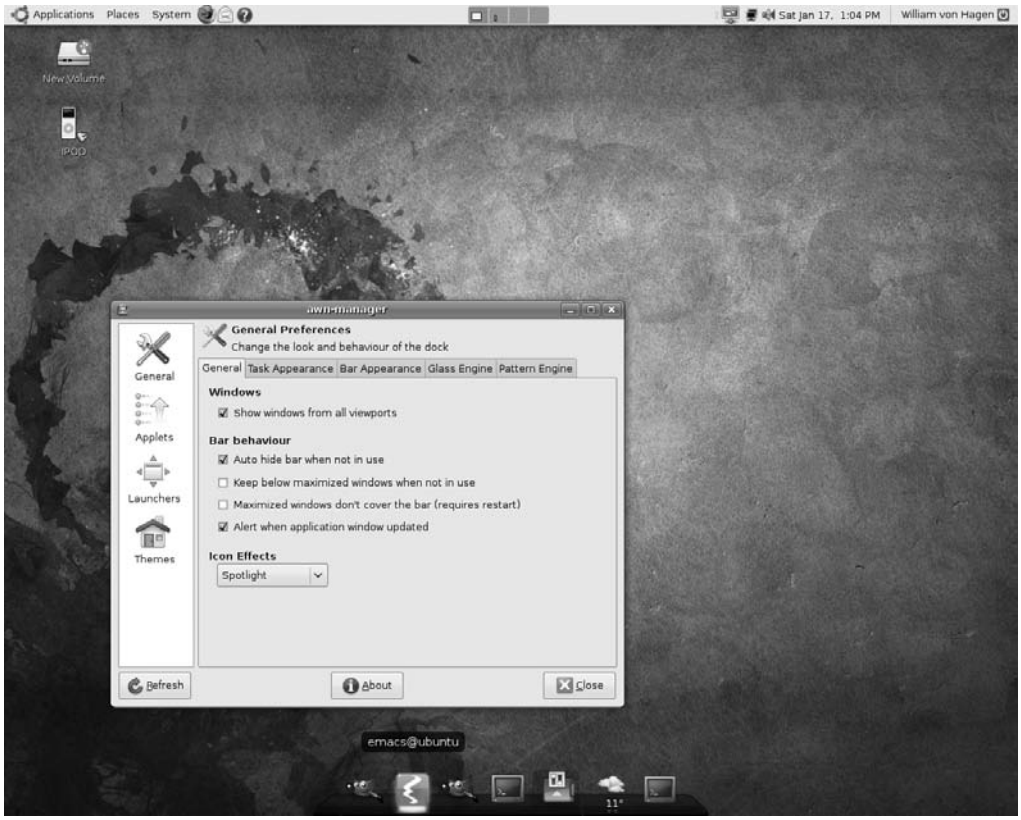
- **General ➪ Auto hide bar when not in use:** This one is actually hard to show because it causes the dock to disappear below the bottom of the screen when it is not in use. To redisplay the dock, move the cursor to the bottom of the screen.
- **General ➪ Alert when application window updated:** This causes the associated icon in the dock to flash whenever a minimized window that is associated with an application is updated. This can be useful if you are executing long-running commands in a minimized window and want to know when they complete.
- **General ➪ Icon Effects ➪ Spotlight:** Selecting Spotlight as the icon effect causes any icon to flash when you position the cursor over that icon in the dock
- **General ➪ Bar Appearance ➪ Look - 3D Look:** Instead of the standard 2D dock background, activating this option causes AWN to use the tapering background shown in Figure 6-15. Compare this against the default flat bar shown in Figure 6-13.

By default, AWN displays only icons for running applications. Two of the nicest features of AWN are the capability to install applets that perform specific functions, and the ability to install new themes to customize the appearance of the dock itself. A good selection of AWN applets is provided with Ubuntu, but you can see the list of available applets (and learn how to build and install them yourself) at `http://wiki.awn-project.org/Applets` and `http://wiki.awn-project.org/Awn-extras`.

TIP **AWN plug-ins enable you to improve the interaction between certain application and their icons in AWN. Plug-ins for audio applications such as Amarok, Audacious, Banshee, and Rhythmbox, and for online chat and instant messaging applications such as Pidgin and XChat, are available at** `http://wiki.awn-project.org/Plugins`.

FIGURE 6-15

The Avant Window Navigator after customization



AWN themes, which enable you to perform additional configuration of the dock's appearance, are much easier to install and use. A large number of custom themes can be downloaded from `http://wiki.awn-project.org/Themes`. One of my favorites, Transparent 3D, is not available there, but you can download it from `www.queervisions.com/arch/2007/10/awn_avantwindow.html`.

Once you download an AWN theme, you install it by selecting the Awn Manager's Themes tab, clicking Add, and navigating to the location where your themes are stored. (I tend to keep them in `~/ART/AWN`, so that I know where to look for them if I want to use them on other systems.) You then select the compressed tar archive file that contains your theme and click Open. Once a theme has been added, the Themes tab looks like that shown in Figure 6-16.

You can then select the theme you want to use and click Apply to have it take effect. Figure 6-17 shows AWN with the Transparent 3D theme enabled.

FIGURE 6-16

The Awn Manager after adding a theme



FIGURE 6-17

The Avant Window Navigator with the Transparent 3D theme



As AWN (and Compiz itself) becomes more popular, more and more applets, application launchers, and themes should become available. Look forward to hours of customization fun. For additional information about AWN, see the Avant Window Navigator Wiki at `http://wiki.awn-project.org`.

TIP **I've been using AWN for quite a while now, and have seen it crash a few times, at which point all of its icons disappear, although the dock itself is still visible. To cause AWN to redraw all of its icons, simply start any other graphical application. All of your icons will reappear!**

Using Emerald and the Emerald Theme Manager

The Emerald window decorator, discussed earlier in this chapter in the section "Compiz History," was originally designed for use with the Beryl compositing window manager. Once the Beryl and Compiz folks began working together, the Compiz Fusion project was formed to port many of the cool features and plug-ins that were created for Beryl so that they would work with Compiz. Luckily for all of us, one item they opted to port and update was Emerald, the window decorator used by Beryl.

As discussed earlier, compositing window managers such as Compiz and Beryl focus on standard window management tasks and compositing, using a separate application to handle window decoration tasks. This division of labor provides a flexible, easily extended desktop environment by enabling different window decorators to be used with a single compositing window manager. By default, GNOME-based Ubuntu systems running Compiz use the `gtk-window-decorator` window decorator, and Kubuntu (and other KDE) systems running Compiz use `kde-window-decorator`. The purpose of these window decorators is to re-implement the look and feel of the traditional GNOME (Metacity) and KDE (KWin) desktops and default window managers in the Compiz environment. Emerald provides an alternative to these that works in either desktop environment. It provides a substantially richer graphical desktop environment because its focus is on creating something new, not preserving a traditional look.

Emerald is not installed by default as a part of Compiz, but you can easily install it using any of the package management utilities that are discussed in Chapter 23. Installing Emerald will also install the library that it uses and the Emerald Theme Manager, a utility that makes it easy for you to configure general aspects of Emerald, as well as to retrieve, install, and use different Emerald themes. When installing Emerald, you should also install the Subversion source code control systems because this software is used by the Emerald Theme Manager to retrieve Emerald themes from remote theme repositories.

Once Emerald has been installed, before you can configure it you must enable the Compiz Window Decorator plug-in (using CompizConfig) and then start the Emerald window decorator by executing the following command:

```
emerald --replace &
```

This starts Emerald and tells it to replace any window decorator that is currently running. After executing this command, you will notice that the default controls on any open windows on your screen change to use Emerald's default theme. You can now start the Emerald Theme Manager by selecting System ➪ Preferences ➪ Emerald Theme Manager. The Emerald Theme Manager is shown in Figure 6-18, which also shows that the program thinks of itself as the Emerald Themer, which is how I'll refer to it from this point on.

FIGURE 6-18

The Emerald Theme Manager, aka the Emerald Themer



As you can see, the initial dialog displayed by the Emerald Themer isn't very interesting, primarily because no themes are provided by default. To download and install default sets of themes, use your favorite Web browser to visit sites such as the following and look for themes labeled "Beryl Emerald Theme":

- `www.compiz-themes.org/`
- `http://compizfusion.blogspot.com/2008/01/compiz-fusion-emerald-themes.html`

You may need to click on the name of a theme that you're interested in so that you can see its Download button. Once you've downloaded any themes that you're interested in, you can load them into the Emerald Themer by clicking Import and navigating to the location where you downloaded the theme(s). Emerald themes typically have the `.emerald` file extension. After selecting a `.emerald` file, click Open to add that theme to the list in the Emerald Themer. Figure 6-19 shows the Emerald Themer with several new themes loaded, but still using its default theme.

Once some number of themes have been retrieved and installed, click any theme to begin using it, as shown in Figure 6-20, which should look somewhat familiar to Mac OS X users.

Experimenting with themes is fun. Selecting a theme that appeals to you can help make your computer system reflect your personality and aesthetics. My personal favorite this week is SolidSlateModified, and I forget what it was last month.

If none of the default themes that you've downloaded are "just right" and you are artistically inclined, you can even create (and optionally publish) your own themes by following the instructions provided at `http://wiki.compiz-fusion.org/Decorators/Emerald`.

FIGURE 6-19

Available themes in the Emerald Theme Manager

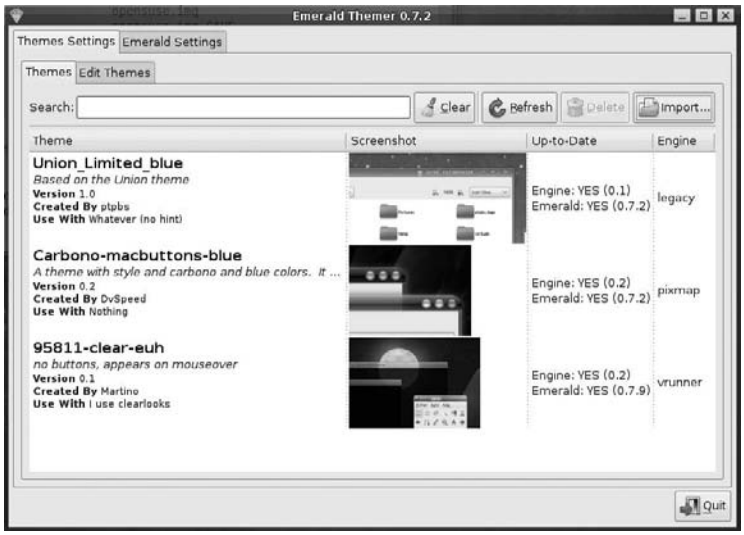
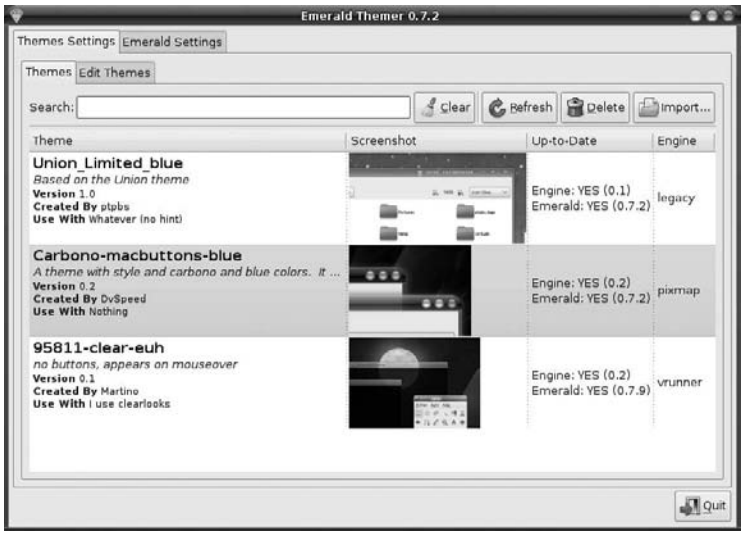


FIGURE 6-20

Using a new theme in the Emerald Theme Manager



Using Compiz with KDE

At present, I generally advise against using Compiz with KDE-based systems such as Kubuntu, primarily because KDE 4's enhancements to the KWin window manager provide much of the desktop functionality that Compiz does. However, if you're using KDE 3, it's possible (and fun) to run KDE 3 with Compiz, as explained in the rest of this section.

Compiz is not installed by default on Kubuntu systems. You can install Compiz by using any of the package management software described in Chapter 23 to install the compiz meta-package. The Compiz package provides the core components of Compiz, a set of standard plug-ins, and a window decorator using the Gtk toolkit. You will therefore also want to install the following KDE-related packages:

- `compizconfig-settings-manager`: The Compiz configuration utility discussed earlier in this chapter, this utility is as useful under KDE as it is under GNOME. The `desktop-effects-kde` application does a good job of enabling basic effects on a Kubuntu (or other KDE) system, but CompizConfig gives you much more granular control over your settings.
- `compiz-kde`: This package installs an appearance preferences tool for KDE.
- `desktop-effects-kde`: This package installs a KDE-based Compiz setup tool.
- `kicker-compiz`: This package installs a modified kicker applet for use with the viewports model used by Compiz.
- `kicker-taskbar-compiz`: This package installs a modified kicker taskbar applet for use with Compiz and the kicker.

You will probably not want to install the `compiz-gnome` package; most people won't want to use this window decorator to make their KDE environment look like GNOME. You can also install Emerald as described earlier in this chapter, but that's completely up to you.

Once all of the software has been installed on your Kubuntu system, using Compiz with KDE simply requires that you start the Compiz window manager rather than KWin when KDE 3 starts. You can do this by setting the environment variable KDEWM (in your `~/.login`, `~/.profile`, or `~/.bash_profile` file) to `compiz` (no path necessary). This entry should look like the following:

```
export KDEWM=compiz
```

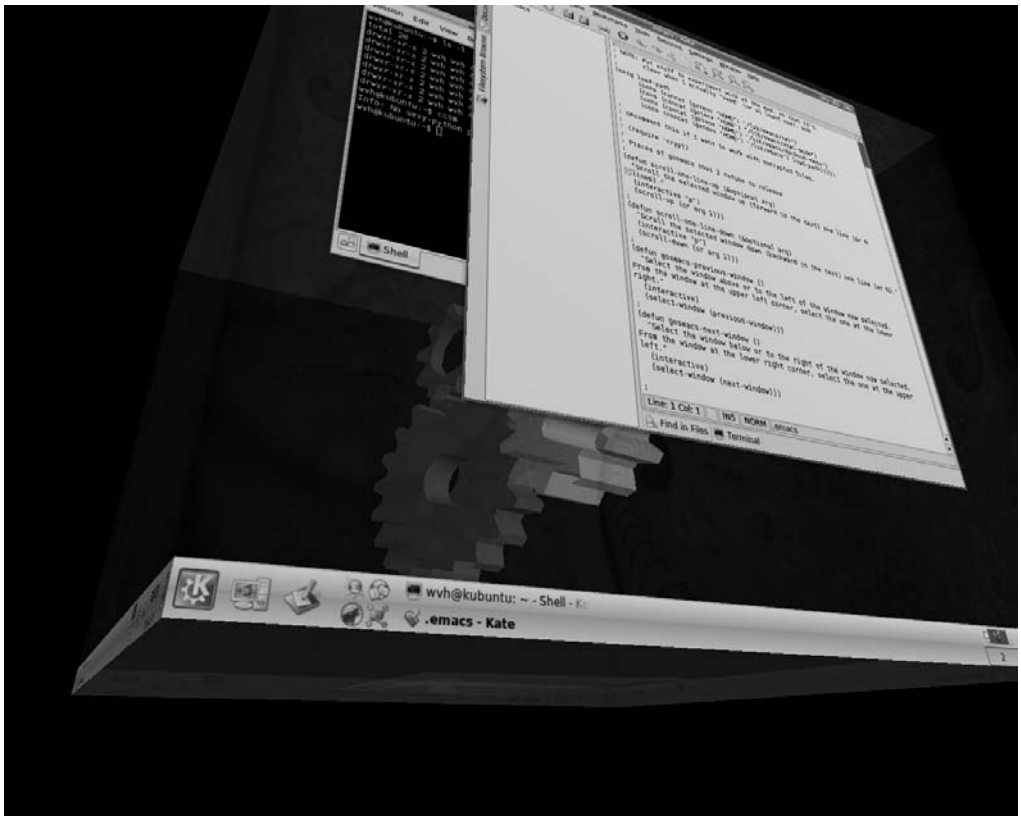
Figure 6-21 shows KDE 3.5 running Compiz as its window manager. To verify that this is indeed KDE and Compiz, this screenshot shows a transparent cube being rotated, with the Cube Gears plug-in running.

If this doesn't work for you for some reason, an alternative approach is to use the KDE Display Manager to start Compiz, instead of KDE, when you log in. You can then get access to most of your traditional KDE applications simply by starting the kicker from any command-line environment.

Regardless of how much I like Compiz, using Compiz with KDE is probably a short-term solution for KDE fans. KDE 4's desktop comes with an impressive suite of compositing features and widgets that should provide the same endless customization and tweaking opportunities that Compiz does for GNOME fans.

FIGURE 6-21

KDE 3.5 with Compiz



Summary

Compositing window managers, specifically the Compiz window manager, have been the darlings of the Linux world for a while now. Compiz is now the default window manager installed and started on GNOME-based Ubuntu systems, and the upcoming release of KDE 4 promises the same level of out-of-the-box compositing support for Kubuntu fans. This chapter explained the basic concepts behind compositing window managers and provided a quick history of compositing on Linux and of Compiz itself. The chapter then explained how to configure Compiz, locate and install plug-ins to improve its capabilities, and install and use various Compiz-related utilities such as the CompizConfig Settings Manager, the Avant Window Navigator, and the Emerald window decorator. The chapter concluded with a discussion of how to install and use Compiz on Kubuntu systems that are running KDE 3.

Chapter 7

Using the KDE 3 Desktop

As introduced in the section of the previous chapter entitled "What's a Desktop? Graphical Environments for Linux," KDE (the K Desktop Environment, `www.kde.org`) is an open source operating environment designed to provide typical computer users with a standard, easy-to-use graphical desktop interface. As with GNOME, the primary goal of KDE is to provide users with a powerful, flexible, and user-friendly graphical environment that supports standard desktop capabilities such as dragging and dropping; provides a consistent look and feel for a wide range of applications; provides a development framework that simplifies developing new graphical applications; and is configurable so that you can customize the desktop environment to work the way that you want it to.

Conceptually, KDE is very similar to GNOME because both are desktop environments, but KDE focuses on providing a rich feature set with extensive customizability, while GNOME focuses on simplicity. The goal of both is usability, but their approaches to usability are very different. In the desktop world, selecting between GNOME and KDE is a Jaguar/Mercedes/BMW decision—which high-quality alternative best suits your needs?

This chapter describes KDE 3.5.*x*, which is the version of KDE that is used on Kubuntu 8.04 and earlier systems. Kubuntu 8.10 and later systems use KDE 4, which although quite beautiful and certainly the future of KDE, is actively under development and very much a work in progress. Both the KDE 4 versions of applications that existed in KDE 3 and new, KDE 4-only applications are still actively changing. Therefore, most of the discussion of the KDE applications in this book focuses on the KDE 3 versions of those applications. If you're using Kubuntu 8.10 or later, you'll want to start with Chapter 8 instead of this chapter. If you're still making your mind up about whether to use Ubuntu, a KDE 3-based Kubuntu, or a KDE 4-based Kubuntu, see the section of Chapter 2 entitled "Ubuntu, Kubuntu, and Which Kubuntu" for more information.

IN THIS CHAPTER

- Using the mouse
- Learning KDE basics
- Using the desktop
- Personalizing the desktop
- Using the Dolphin file manager
- Using the Konqueror browser
- Using a window manager
- Installing KDE on GNOME systems

NOTE **KDE can easily be installed on an Ubuntu system, just as GNOME can easily be installed on a Kubuntu system. This level of choice and flexibility is one of the greatest things about open source software. Therefore, whether you've started with GNOME or KDE, you can easily install the other on your system (without deleting or disabling the current desktop environment), and switch between them at your leisure until you find the desktop environment that is right for you. If you're an Ubuntu user who is curious about KDE, see the section "Installing KDE on a Ubuntu System" near the end of this chapter for information about installing and starting KDE on a standard, GNOME-based Ubuntu system.**

A Short Course in KDE History

KDE was initially inspired by CDE (the Common Desktop Environment), an X Window system-based graphical environment that was cooperatively designed by major UNIX vendors such as IBM and Sun Microsystems and originally ran on those vendors' hardware. CDE was designed and built using the Open Group's Motif X Window system libraries, which meant that any vendor wishing to deploy CDE had to license Motif. This was unacceptable in the free, open source environment of Linux, so KDE was initially written to provide a free alternative to CDE. The KDE project (originally, KDE stood for Kool Desktop Environment) was proposed by Matthias Ettrich in Usenet posts on October 14, 1996, making it the first desktop environment developed for Linux. Over the years, the KDE project has been supported by many Linux vendors and fans, including Caldera, Corel, Delix, Intel, Novell, Mark Shuttleworth (Mr. Ubuntu himself), O'Reilly Associates, and SUSE.

KDE might have become the one true desktop environment for Linux and other open source systems were it not for two factors:

- The KDE folks selected a widget and windowing package called Qt to serve as its graphical underpinnings, which initially had some licensing issues.
- One of the unwritten rules of open source software is that anything worth doing is worth doing several times.

I'm just kidding about the second point (although it often seems to be the case), but the first one is completely accurate. Qt is a cross-platform application framework that is a product of a company called Trolltech (`www.trolltech.com`). A free version of Qt has always existed, but was licensed with provisions that modified versions could not be redistributed, which violated the key "redistributable" provision of the Free Software Foundation's definition of open source software. Similarly, members of the open source community were concerned because there was no legal commitment from Trolltech to continue to provide a free version, and there was no guarantee that the free version of Qt would remain free if Trolltech were acquired by a company that was disinclined to offer it.

These issues were resolved in 1998 with an agreement between Trolltech and KDE project members which states that Trolltech will release a new free version of Qt at least every 12 months and, should this release not occur, the last free version of Qt would automatically be re-licensed under the BSD free software license. This license has not changed with Nokia's acquisition of Trolltech in early 2008—if anything, the Qt licensing has gotten even better because Qt now uses a GPL version 3 license (`www.gnu.org/licenses/gpl-3.0.html`).

Although the Qt licensing issues were resolved long ago, a side effect of these issues was the formation of the GNOME project, which many Linux distributions (such as Fedora, Ubuntu, Red Hat, Yellow Dog, and so on) and their devotees have come to prefer over KDE as their default desktop environment. The key aspects

of using the GNOME interface were discussed in the Chapter 5, and both GNOME and KDE applications are discussed throughout this book. The remainder of this chapter discusses the basic capabilities, characteristics, and use of the KDE 3.*x* desktop. Either KDE 3 or KDE 4 is available as the default desktop environment for Linux distributions such as Kubuntu, Mandriva, MEPIS, OpenSUSE, Red Flag, and so on, and the "Which is better, GNOME or KDE?" discussions on the Internet have wasted almost as much time and space as the "vi vs. Emacs," "tastes great/less filling," and other "religious" discussions.

Using the Mouse

Interacting with a graphical user interface (GUI) of any kind traditionally requires the use of a pointing device (a mouse with a desktop computer and a "trackpad" of some kind with a laptop computer) and keyboard to control the cursor to select, move, and click items; access menus; and enter data within applications. As you'll see throughout this book, the actions associated with the different mouse buttons do different things depending on the portion of the screen or application in which you use them. This is known as being context sensitive, and is one of the keys to performing many different types of actions using only two or three little buttons.

NOTE **The default mouse actions in KDE are different from those in GNOME. For example, left single-clicking a desktop icon in KDE opens the object associated with that icon, while in GNOME, the icon is simply selected. If you are familiar with GNOME but new to KDE, this can take some getting used to. Or, as I do, you may simply want to change it back to the standard double-click model. See the section later in this chapter entitled "Customizing Mouse Behavior" for information on changing this.**

By default, the mouse in KDE is configured for a right-handed user, with the different mouse buttons assigned to the following actions:

- **Left button:** Single-clicking an icon, menu, menu item, or object within an application "executes" or opens the item associated with that icon. Holding down the left button and dragging the cursor in an application or on the desktop enables you to select multiple icons or objects within that application. Holding down the left button over an icon and dragging that icon moves that object. These operations can be (and often are) combined. For example, in a word processor, you frequently left-click once to position the cursor, hold down the left button and drag the cursor to select a block of text, and then left-click, hold down the mouse button, and drag the selected text to a new location in your text.
- **Middle button:** Single-clicking on the desktop displays the Window List menu, which enables you to select and display a specific window, and to unclutter or cascade windows. The latter two options reposition the active windows on your screen to make it easier to find a specific window. Single-clicking an icon on the desktop or in an application opens the item associated with that icon, something like the Expose function on Mac OS X.
- **Right button:** Displays any context-sensitive menu that is available at the current cursor position.

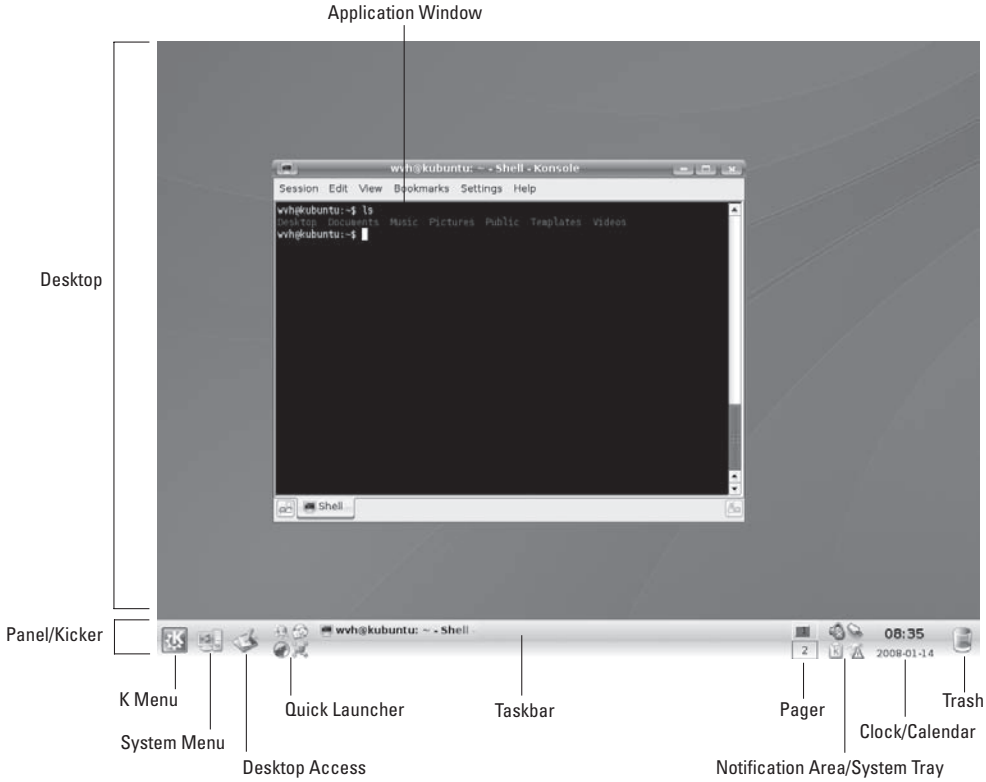
You can configure how responsive the mouse is, what your mouse pointer looks like, and what the various mouse buttons do by selecting System Settings ➪ Keyboard and Mouse ➪ Mouse configuration tool, which is explained later in this chapter in the section "Customizing Mouse Behavior." For example, if you are not right-handed and are fortunate enough not to have been forced to use the mouse with your right hand, you can configure the mouse for left-handed use (which reverses the meaning of the various buttons) using this application. Other aspects of mouse behavior, such as the actions that different mouse buttons perform on the desktop, are configured by selecting System Settings ➪ Desktop ➪ Behavior tool.

KDE 3 Desktop Overview

Figure 7-1 shows the default KDE desktop on a Kubuntu 8.04 or earlier desktop system the first time you log in. This figure also displays a single application window as an example—you won't see that unless you explicitly select the K Menu ➪ System ➪ Konsole - Terminal Application menu option, but I wanted to be able to explain the window controls that are available in any KDE application window.

FIGURE 7-1

The default KDE 3 desktop on Kubuntu 8.04 and earlier



The KDE folks and the Kubuntu team have gone to a lot of effort to create an attractive, eye-pleasing background, a set of fonts, window decorations, and control buttons for any applications that you start on your Kubuntu system. These window decorations and controls are known collectively as a theme, and you can easily customize it yourself, as explained later in this chapter.

Figure 7-1 identifies the various portions of the default KDE 3 desktop. Aside from the workspace where application windows are displayed, the most significant portion of the Kubuntu desktop is the panel, commonly referred to as the *kicker*, which is the equivalent of the Microsoft Windows task bar or the Mac

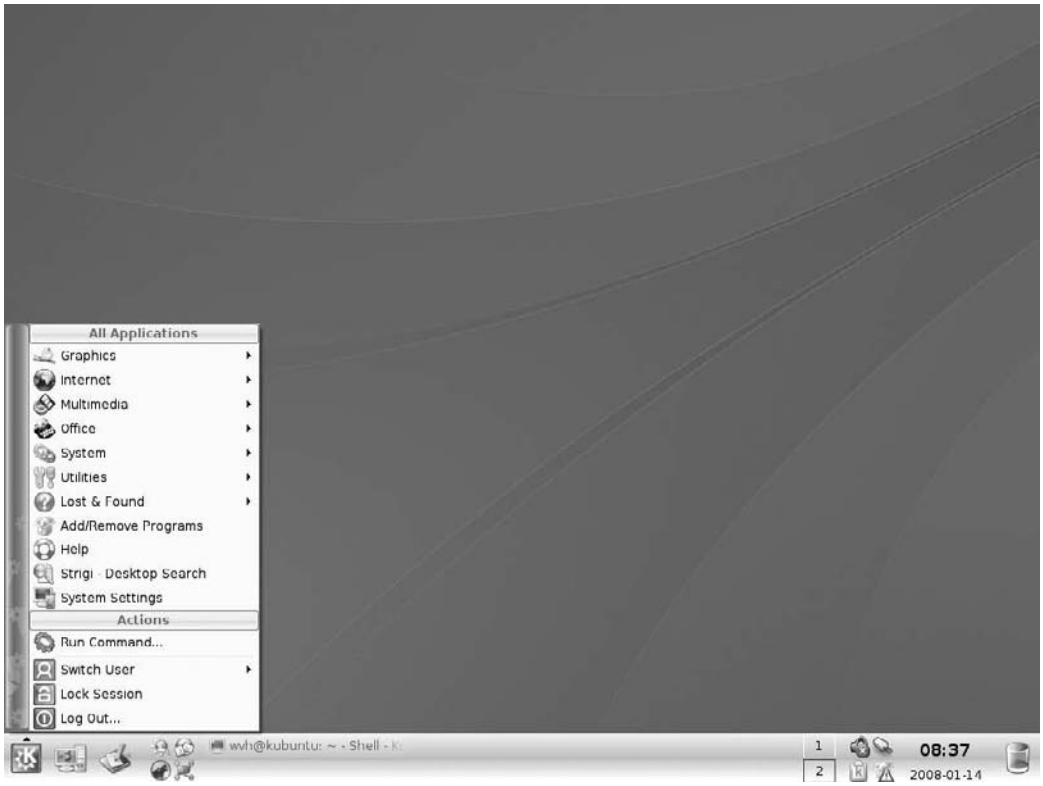
OS X dock. A panel is a special portion of your desktop, controlled by KED's kicker application, which provides access to various ways of interacting with the desktop and launching different types of applications. KDE can display panels along any edge of the screen, but the default KDE configuration displays a single panel across the bottom of the screen. The panel contains the icons and work areas that are associated with applets, which are lightweight applications that are run in the context of the panel and provide various capabilities. These are known as applets because they are not standalone applications, but require both the panel and KDE infrastructure in order to function correctly. See "Customizing Panel Contents," later in this chapter, for more complete information about any of these applets.

The applets in the default Kubuntu panel are as follows:

- **K Menu:** The K Menu is a special applet that provides the primary menu for applications and actions on the KDE desktop. A menu consists of a hierarchical list of shortcuts to specific commands, tasks, and applications, and can also contain other menus (known as "submenus"). To open the K Menu, click once on the K icon in the lower-left corner of the screen. Figure 7-2 shows the default K Menu.

FIGURE 7-2

The default K Menu on KDE 3-based systems



The contents of the K Menu provide easy, graphical access to applications, computer system hardware, and system-related tasks, and are described in detail in the next section, "Menus in KDE." To display any sub-menu, you simply left-click on the name of the menu, and the associated menu drops down and displays on your screen. To close a menu that is being displayed, simply left-click elsewhere on the screen outside of the perimeter of that menu.

- **System Menu:** A custom instance of a generic applet that enables you to add a menu to a panel, the System Menu applet provides one-click access to physical or logical locations on your system. It does this by starting the Dolphin file manager to display locations such as your Documents folder, your Home folder, system Storage Media, Remote Locations (on the network), and the Users folder on your system.
- **Desktop Access:** Hides or exposes all applications that are currently running on your desktop.
- **Quicklauncher:** provides one-click access to commonly-used applications such as the Amarok audio player, the Konqueror web browser, the Kontact personal information manager, and the Kopete instant messaging client.
- **Taskbar:** The Taskbar is a special panel applet in which each application that is active on your current workspace displays an associated control region in the central portion of the KDE panel. Right-clicking on this control minimizes and maximizes the application, while left-clicking on this control displays a context-sensitive menu for moving and controlling that application's window(s).
- **Pager:** An applet that enables you to manage and switch between multiple workspaces and provides a miniature display of each active workspace.
- **Notification Area/System Tray:** An applet that provides a display area, much like that found on Microsoft Windows systems, in which applications and applets can display icons that represent system messages, such as receiving e-mail or low-power conditions, or provide access to system resources such as the clipboard used for cutting and pasting data between applications. One important application that displays messages here is the Adept Updater, which displays an icon in the notification area to let you know that system updates are available.
- **Clock/Calendar:** An applet that enables you to view and set the system time, time zone, and date.
- **Trash:** An applet that provides access to the Trash Can, which is a special portion of the KDE desktop to which you can drag files to subsequently delete them.

As you'll see throughout the rest of this chapter, KDE is extremely configurable—the chances are that no two users' KDE desktops will look the same, which is completely fine. How you configure and use your computer system is up to you, and should reflect your interests and the types of things that you want to do. The job of desktops such as KDE is to make it as easy and as comfortable as possible to accomplish whatever task you are trying to do.

KDE 3 Application Windows

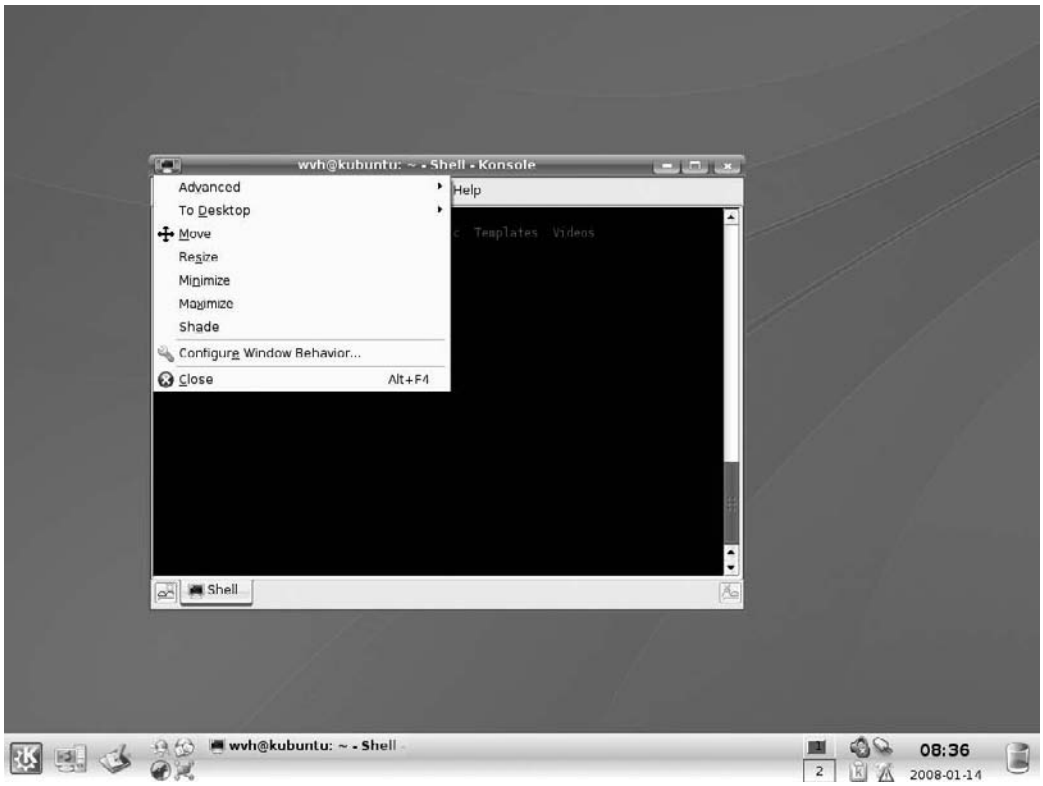
All KDE 3 windows, such as the Konsole application window shown in Figures 7-1 and 7-2, provide a number of window controls that enable you to move the window around on the screen, control its size, and so on. The window manager that you are using determines the location of these controls, the number

of controls displayed, and the capabilities provided by these controls. These window controls are shown in the top part of the frame around the application window in most window managers, a feature typically known as the title bar.

The control at the far left of the title bar displays a menu of window control commands. The contents of this menu depend on the window manager that you are using. The window control menu displayed by the default window manager used by KDE on Kubuntu systems, KWin, is shown in Figure 7-3.

FIGURE 7-3

The KWin window manager's window control menu

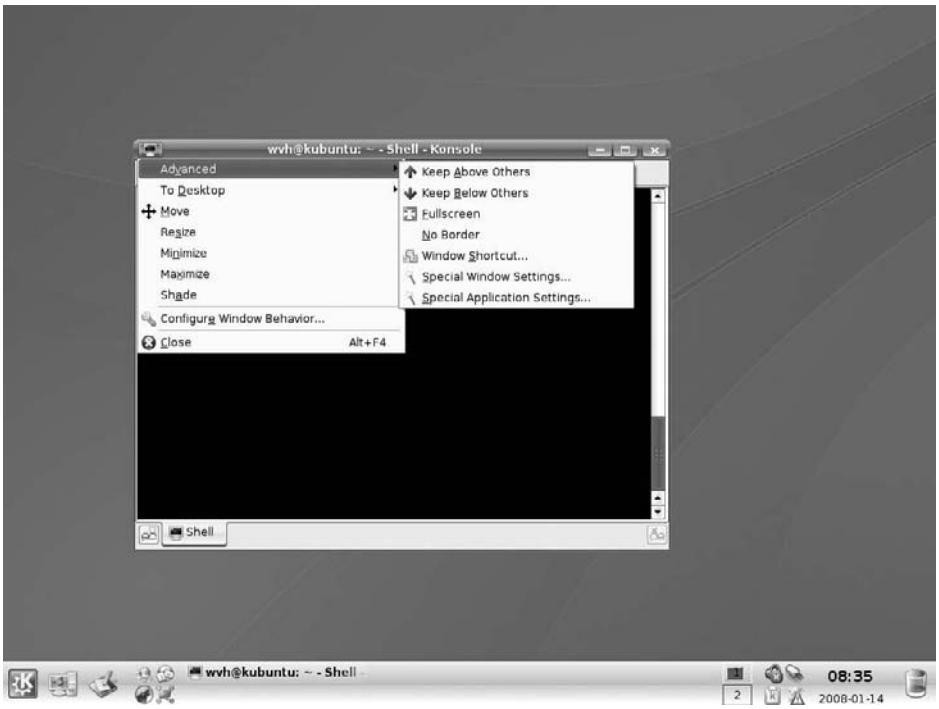


The action associated with each of these menu options is fairly straightforward, with the exception of the following:

- **Advanced:** Displays the secondary menu shown in Figure 7-4, which provides additional controls over the behavior of the current window.

FIGURE 7-4

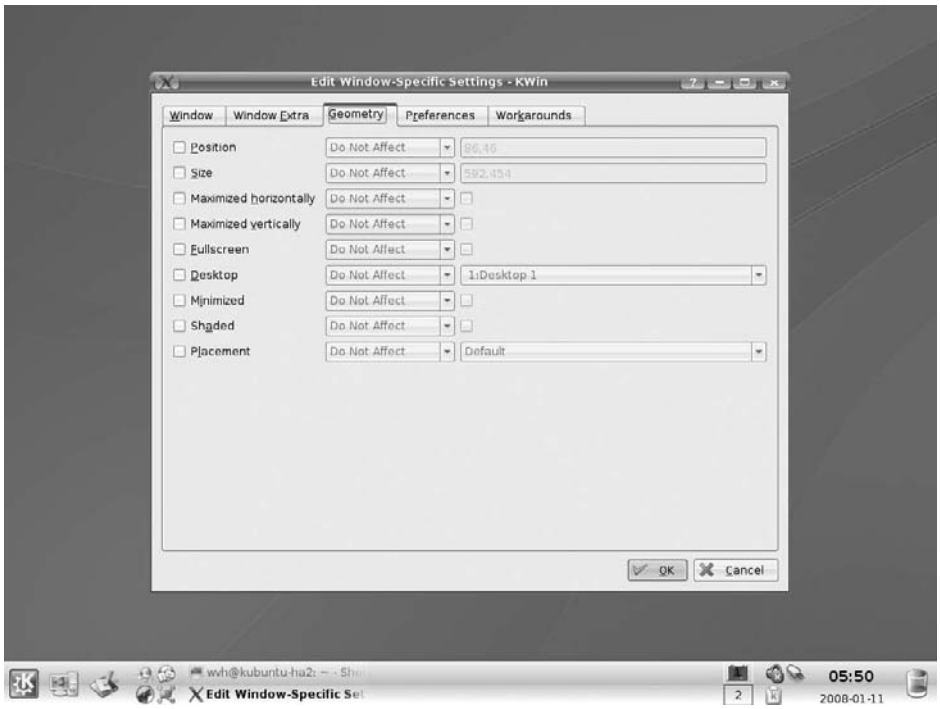
The KWin window manager's advanced window control menu



- **Keep Above Others** and **Keep Below Others:** Enables you to make sure that the current window is always visible (Above) regardless of how many windows are on the screen, or that the current window is always displayed under all other open windows (Below). The Fullscreen command expands the current window to use the entire screen—this differs from Maximizing the window in that full-screen windows no longer display a title bar, and must be reduced in size by using the Alt+F3 key combination to re-display the Window Control menu, and re-selecting the Advanced menu's Fullscreen option. The Special Window Settings and Special Application Settings menu options display dialogs like the one shown in Figure 7-5 that enable you to configure the behavior of the current window or all other windows for the same class of application as the current window.
- **Shade:** Shrinks or "rolls up" the current window so that only its title bar is shown. This can be convenient if you want to keep a large number of windows on the screen without minimizing them to the Taskbar section of the panel.

FIGURE 7-5

The advanced window control menu's window-specific settings dialog



- **Configure Window Behavior:** Displays the dialog shown in Figure 7-6, which enables you to configure general aspects of window behavior on your Kubuntu system. This dialog is essentially the same as the Window Behavior configuration item in KDE's System Settings configuration application. The System Settings application is discussed later in this chapter in the section "Customizing Your Desktop."

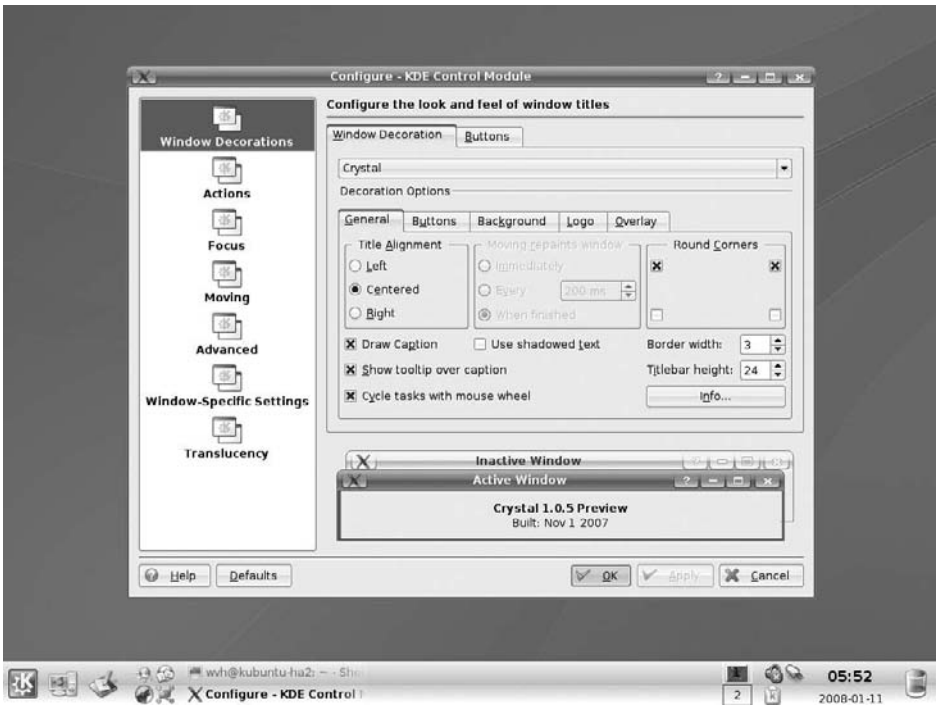
The three window controls at the right of the title bar in Figure 7-3 do the following, from left to right:

- Minimize the current window.
- Switch between a full-screen or a user-specified window size.
- Close the current window, terminating any processes that it is running.

All of these window controls are fairly standard, and are found in the window decoration and title bar layouts that are provided by most modern window managers.

FIGURE 7-6

The window control menu's configure Window Behavior dialog



Menus in KDE 3

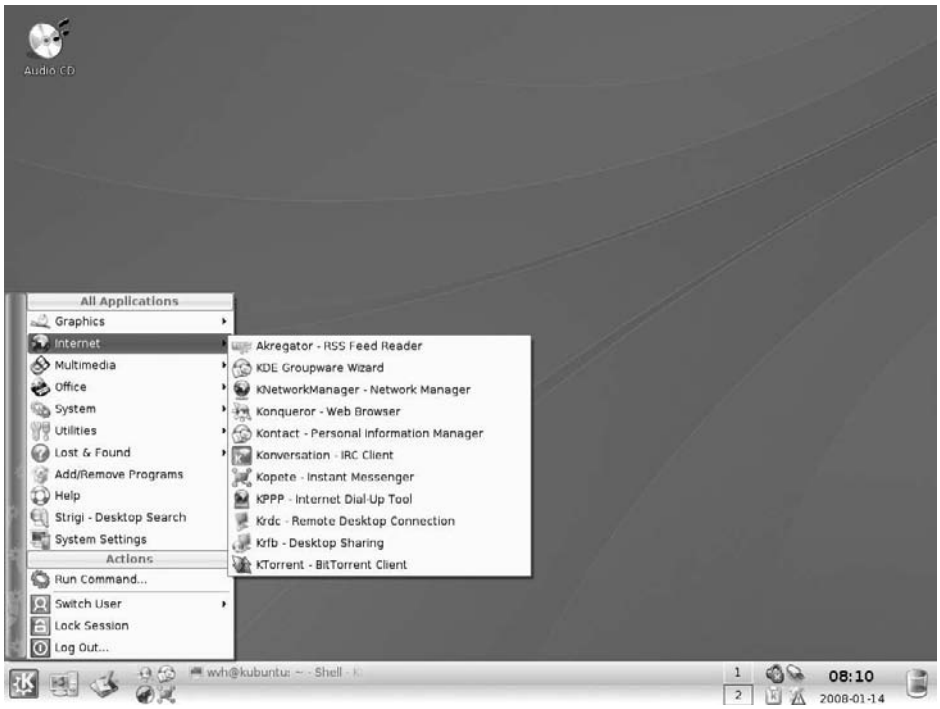
As discussed earlier in this chapter, KDE provides two different types of menus: menus located in a panel, such as the K and System menus, which provide easy access to different applications, locations, and capabilities on your system; and context-sensitive menus that provide access to different capabilities depending on the context (application, desktop, and so on) in which where you've activated them. This section provides a more detailed discussion of these different types of menus—including how to display them, the types of actions that they provide access to, and how to customize them.

Panel Menus

Menus are displayed in the panel when you left-click on the K Menu icon, System Menu icon, or the icon for any other menu-related applet that you have added to the panel. Menus either contain entries for specific tasks or locations on your system, or contain other menus. Menus located on other menus are known as submenus. For example, Figure 7-7 shows the default contents of the K Menu's Internet submenu, which I discuss in more detail in the next section.

FIGURE 7-7

Menus and submenus in KDE



The K Menu

The K Menu makes it easy to launch different types of KDE applications, grouping them into submenus that display when you move the cursor over each menu entry. The K Menu provides the following groups:

- **Games:** Provides access to games organized into various categories: Arcade, Board Games, Card Games, Games for Kids, and Tactics & Strategy games. For more information about games on Kubuntu, see Chapter 20, "Would You Like to Play a Game?"
- **Graphics:** Provides access to several applications for creating graphics, capturing and manipulating images off the screen or from a scanner, and viewing PDF files, and a small, lightweight image viewer for previewing existing graphics files. See Chapter 18, "Working with Graphics," for more information about graphics software on Kubuntu Linux systems.
- **Internet:** Contains entries for accessing information sources found on the Internet, such as RSS feeds, the Konqueror browser for the Web (and other objects), the Kontact Personal Information Manager (PIM), desktop sharing, instant messaging and IRC clients, and so on.

- **Multimedia:** Provides easy access to audio and video applications, including applications for configuring sound devices and levels, converting your (officially purchased) audio CDs to online format, playing online audio and video, and sophisticated applications for creating different types of CDs, VCDs, and DVDs. See Chapter 19, "Working with Multimedia," for more information about working with different types of multimedia files and mediums on Kubuntu systems.
- **Office:** Provides access to business applications that are available on your Kubuntu system, including the Kontact Personal Information Manager (PIM); a fax address book and utility; a personal time tracker (useful for the consultants among us!); a wizard for connecting to groupware applications such as eGroupware, Kolab, OpenExchange, and GroupWise; the KNotes application for creating and managing Post It-like notes on your desktop; and the OpenOffice.org Drawing, Presentation, Spreadsheet, and Word Processor. See Chapter 18 for more information about the OpenOffice.org Draw package. See Chapter 17 for information about the OpenOffice.org Presentation (Impress) and Spreadsheet (Calc) packages. See Chapter 16 for information about the OpenOffice.org Word Processor (Writer).
- **Settings:** Provides access to a small number of system and service configuration utilities that provide their own configuration framework rather than using the standard KDE/Kubuntu System Settings framework.
- **System:** Provides access to system maintenance, and diagnostic applications, including the Adept Package Manager, the D3lphin (Dolphin) file manager, the KCron Task Scheduler, the KBackup Backup System, the Kubuntu Device Database, performance monitoring and log viewing applications, and—probably most important for many of us—the Konsole Terminal Program.
- **Utilities:** Provides a general category for applications including an archiving tool, a calculator, the Kate text editor, a virtual keyboard, and tools for magnifying the screen and configuring input methods on your system, viewing print jobs, and so on.
- **Lost & Found:** Provides a catch-all menu entry for applications that couldn't be added to other KDE menus for one reason or another. This menu displays only if problems were encountered when trying to create other menu entries.

In addition to these menus, the K Menu provides easy access to popular applications such as an Add/Remove Programs entry (which invokes the Adept Package Installer), centralized help for KDE and many KDE applications, the Strigi Desktop Search application, and a System Settings item that provides centralized access to a number of smaller tools that make it easy to configure KDE. The K Menu also provides shortcuts to common actions, such as running a command without requiring that you start the Konsole application, switching between different users on a single Kubuntu system, and locking the screen if you have to leave for a while but don't want to log out; it also provides a Log Out action that enables you to log out, restart your system, or shut down your system entirely.

The System Menu

Like the K Menu, the System Menu is accessed through an icon in the KDE panel. The System Menu provides convenient shortcuts to the following commonly accessed locations on your system:

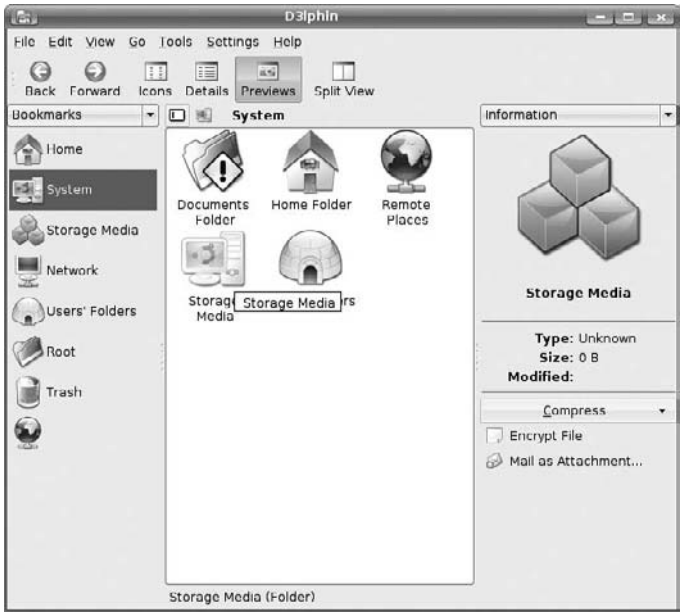
- **Documents folder:** Displays the Documents folder, located in your home directory, in the Dolphin file manager
- **Home folder:** Displays your home directory in the Dolphin file manager

- **Storage media:** Uses the Dolphin file manager to display all of the physical storage media that your system is currently using
- **Remote places:** Uses the Dolphin file manager to enable you to browse all of the network resources that your system is currently aware of
- **Users folder:** Uses the Dolphin file manager to display the home directories for all of the users that have accounts on your system and have directories in the `/home` directory

Figure 7-8 shows the Dolphin file manager providing a higher level view of a sample system. Each item in the System Menu is visible in the left menu and, as an icon, in the central pane.

FIGURE 7-8

The Dolphin file manager



For more information about using Dolphin, see "Using the Dolphin File Manager," later in this chapter.

Context-Sensitive Menus

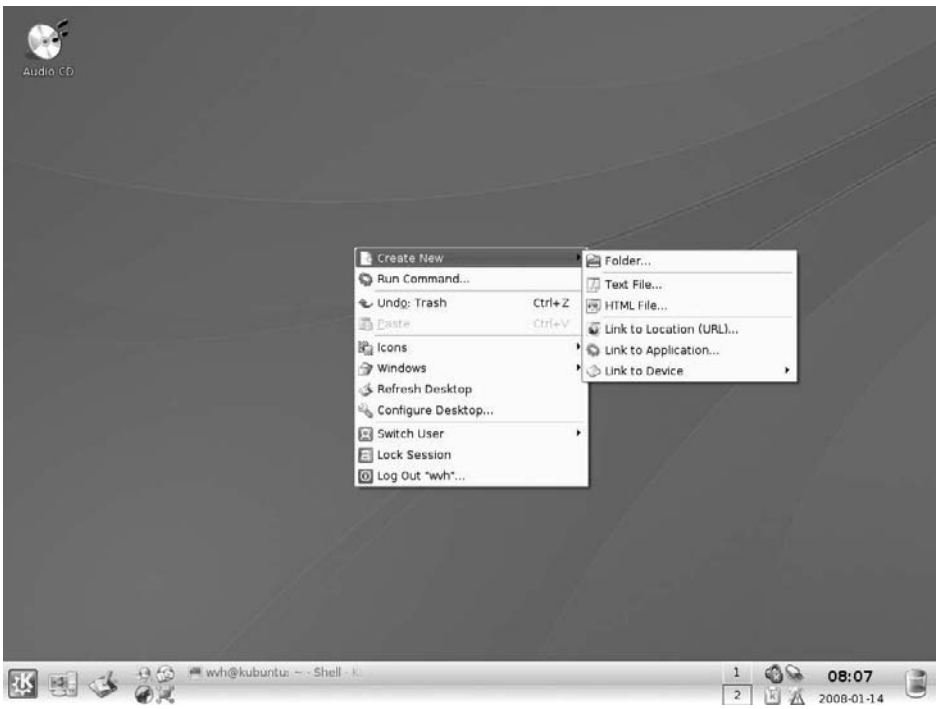
The menus discussed earlier in the previous section always display when you left-click on a specific menu entry in the panel, regardless of what else is happening on your system. Context-sensitive menus display when you right-click on something on your KDE 3-based Kubuntu screen, but their contents depend on what you've right-clicked on. Because their contents therefore differ based on the context in which you've

right-clicked, these right-click menus are referred to as context-sensitive menus. Figure 7-9 shows the context-sensitive menu that you see when you right-click on any portion of the background for your desktop (the larger, left-most menu), and also illustrates that context-sensitive menus can contain submenus.

To demonstrate how application-specific these different context-sensitive menus can be, Figure 7-10 shows the context-sensitive menu that displays when you right-click on the desktop icon for an audio CD. As you can see, this context-sensitive menu displays actions and commands that are specifically relevant to an audio CD. These include options to open the CD with the Konqueror or Dolphin file managers or another application that you specify, copy, or extract the contents of the CD with the k3b application (discussed in Chapter 19), play the CD with the Amarok audio player (also discussed in Chapter 19), and so on. This menu clearly differs from the default menu that displays when right-clicking on the desktop, as shown in Figure 7-9, because different actions and applications are relevant to the desktop and audio CD contexts.

FIGURE 7-9

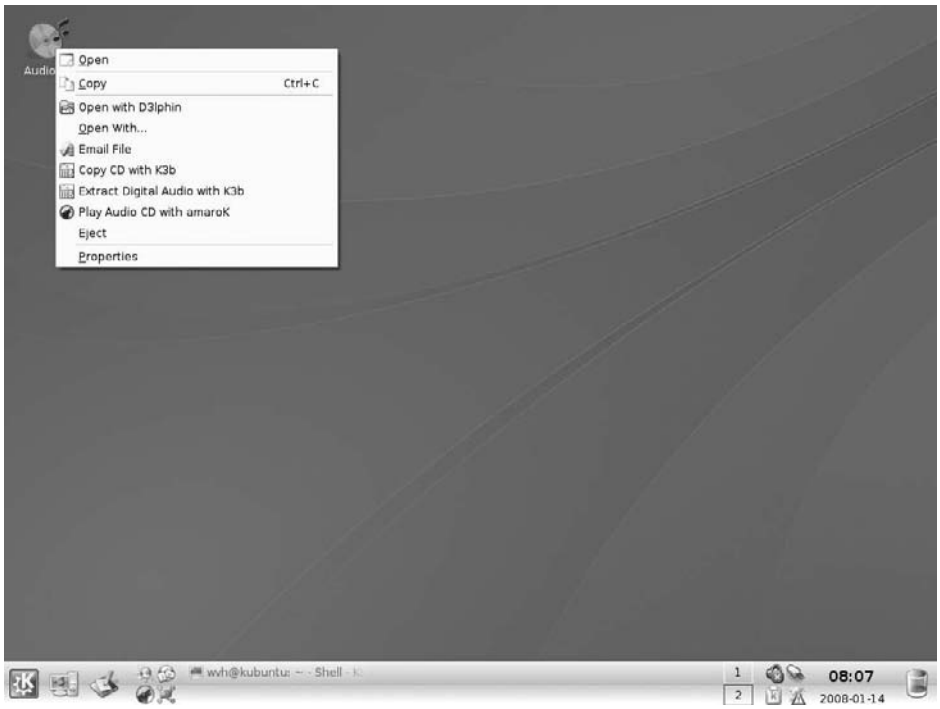
The KDE desktop's context-sensitive menu



The contents of context-sensitive menus are application- or device-specific, which means that you can't directly customize them. In some cases, these menus can be customized by installing additional applications or applets, or (in the case of extensible applications such as the Firefox web browser) by installing extensions or plug-ins that add new capabilities and therefore provide associated menu items.

FIGURE 7-10

The context-sensitive menu for an audio CD



Customizing KDE Menus

KDE 3-based Kubuntu systems provide you with three ways to customize the contents of the K Menu:

- **Customize the K Menu by adding predefined applications:** The K Menu's Add/Remove Programs menu item enables you to customize this menu by installing new software packages from a pre-selected set of popular applications. Using this menu item is discussed in the section of Chapter 23 entitled "Using the Add/Remove Programs Menu Item in Kubuntu." This menu item is extremely convenient to add some popular application to the K Menu, but isn't much use if you want to add your own menus or create custom menu entries for specific applications.
- **Customize the K Menu by adding predefined menu entries and sections:** The panel's Configuration utility enables you to add entries from a pre-selected set of available submenus. See "Configuring the Panel," later in this chapter, for more information.

- **Use the KDE Menu Editor:** To add your own menus or create custom menu entries, you will have to use the KDE Menu Editor. The easiest way to start the KDE Menu Editor is from a context-sensitive menu, by right-clicking on any panel menu icon or on any item on menus such as the K Menu or System Menu.

Figure 7-11 shows the context-sensitive menu that is displayed by right-clicking on the panel's K Menu icon. Figure 7-12 shows the context-sensitive menu that is displayed by right-clicking on any item in a K Menu item that is currently being displayed, in this case the Utilities menu.

Selecting the Menu Editor or Edit Menu contextual menu item displays the KDE Menu Editor, shown in Figure 7-13.

FIGURE 7-11

The context-sensitive menu for the K Menu icon



FIGURE 7-12

The context-sensitive menu for an existing menu item

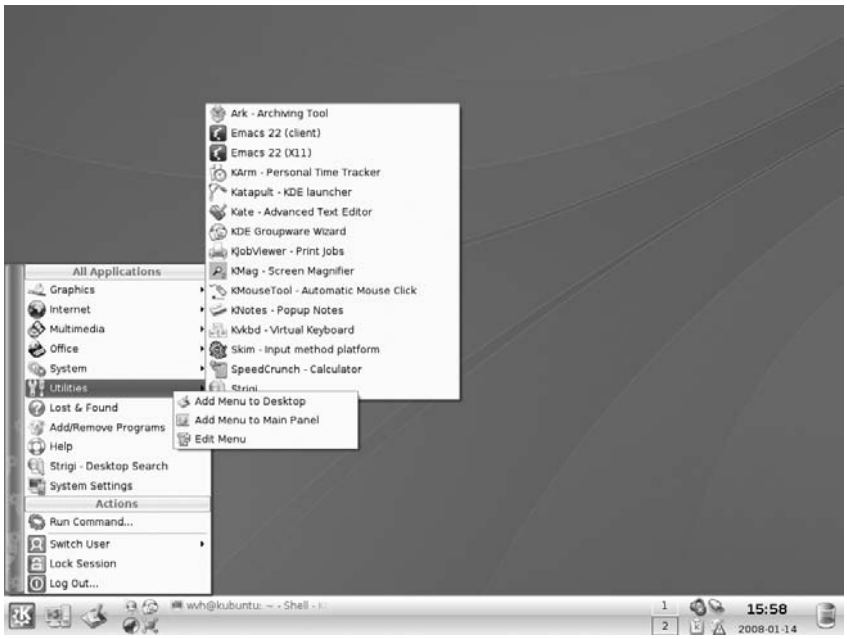
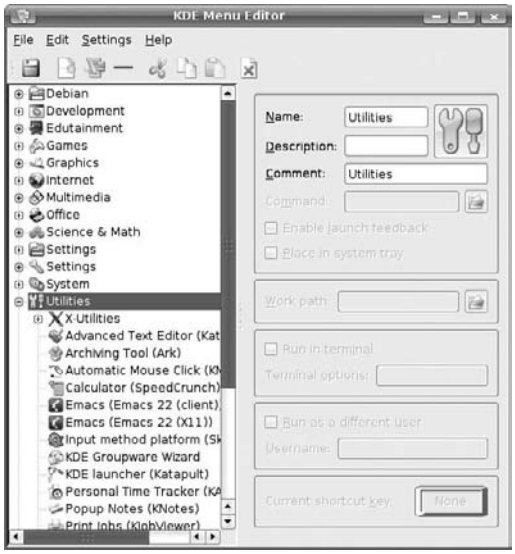


FIGURE 7-13

The KDE Menu Editor



Suppose that you wanted to create your own submenu called Favorites and begin adding entries for your favorite items to it. To do this (using `emacs` as an example of a favorite command), do the following:

1. To add a new top-level menu, select any non-menu item that is at the top level of the K Menu. (To add an item to any of the menus that are already displayed in the window at left, select the menu to which you want to add a menu or submenu.)
2. Select the File ➪ New Submenu option. A dialog like the one shown in Figure 7-14 displays.
3. Enter the name of the new menu that you want to create (I'm using Favorites in this example) and click OK. The new menu displays in the window at left, and the basic information for this menu displays in the right half of the KDE Menu Editor (see Figure 7-15).

FIGURE 7-14

Creating a new menu item

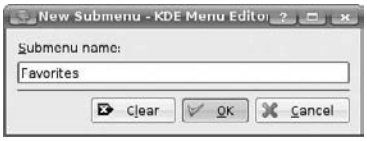


FIGURE 7-15

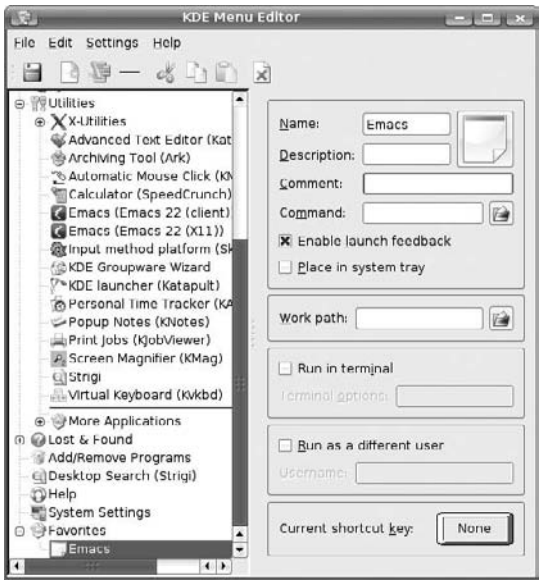
Entering information about a new menu



4. (Optional) Enter a short description and comment for the new menu entry for future reference. You can also change the icon that is displayed to the left of the menu entry by clicking the current icon to display a dialog that enables you to select another icon to display beside the menu heading.
5. To add a new entry to the new menu, make sure that it is still selected in the left pane of the KDE Menu Editor, and select the File ➪ New Item menu entry. A dialog like the one shown in Figure 7-16 displays.
6. Enter the name of the new item that you want to create (I'm using Emacs in this example) and click OK. The new menu entry displays in the window at left, and the basic information for this item displays in the right half of the KDE Menu Editor, as shown in Figure 7-16.

FIGURE 7-16

Entering information about a new menu item

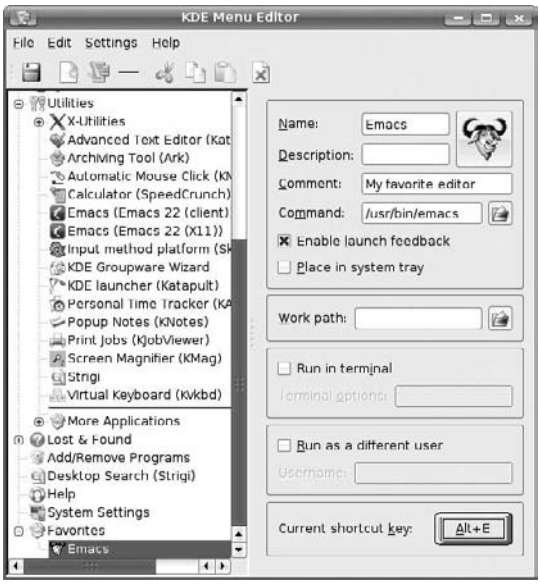


7. Enter the full pathname of the command that you want to execute when this menu item is selected (in this case, `/usr/bin/emacs`). (You can also click the folder icon to the left of the command input area and browse your computer to find the location of the command that you want to associate with the new menu entry.) Optionally, you can enter a Description, Comment, and Work path (the directory that the application will start from); click Run in terminal to automatically start the application in its own terminal window; click Run as a different user (to always run the application as a specific user); and bind a special shortcut key to the new command by

clicking the Current shortcut key button (labeled None by default) and typing the key combination that you want to use to select this menu item and execute the associated command. Figure 7-17 shows the same dialog as that shown in Figure 7-16, but with a Comment, the correct command, a custom icon, and the shortcut key Alt+E bound to this menu item.

FIGURE 7-17

Completed information about a new menu item



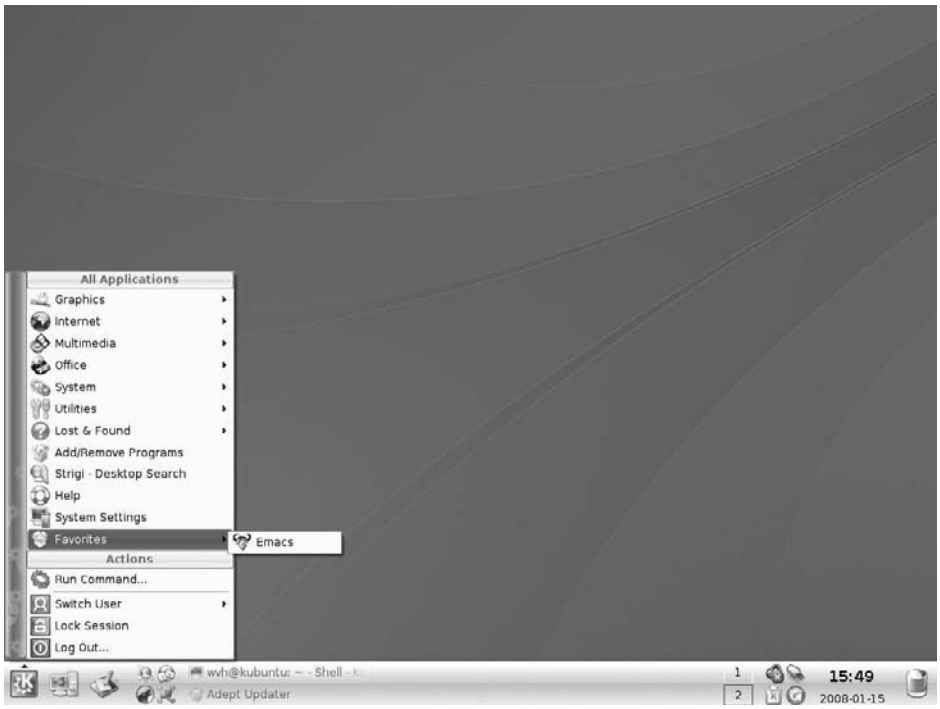
8. To save your changes, select the File ➪ Save menu item or type Ctrl+S. A dialog displays that provides progress information as your changes are saved and the system updates its menu configuration.
9. To exit from the KDE Menu Editor, select the File ➪ Quit command. If you have made changes but have not saved them, a dialog displays that gives you the option of saving them before exiting. Otherwise, the KDE Menu Editor dialog simply closes.

At this point, redisplaying the K Menu will show the new Favorites menu and the Emacs item that it contains, as shown in Figure 7-18.

The flexibility provided by the KDE Menu Editor makes it easy for you to create custom menus and menu entries for locally written or locally installed applications. This can be especially useful in academic or corporate environments, where you want to make sure that users can easily locate and execute sets of site-specific applications.

FIGURE 7-18

The new Favorites submenu and Emacs menu item



Customizing the KDE 3 Desktop

The previous section showed how to customize the menus that provide graphical access to the applications on a KDE 3-based Kubuntu system. Luckily, every part of the KDE desktop is similarly configurable. While the GNOME version of Ubuntu provides a Preferences menu to hold menu entries for the many small applications that configure specific portions of the GNOME interface, KDE's K Menu provides a single menu entry called System Settings that displays an application in which icons for many different configuration items can be selected. The System Settings application is a replacement for the configuration utility known as the "Control Center," which was used in earlier versions of KDE. Selecting the System Settings menu item from the K Menu initially displays the dialog shown in Figure 7-19.

The next few sections highlight how to do some of the most common tweaks that most people do to make their systems easier to use and generally personalize them. Many other customizations are available, but discussing all of them could make this book impossible to lift. Let's focus on the highlights!

FIGURE 7-19

KDE 3's System Settings application



Customizing Mouse Behavior

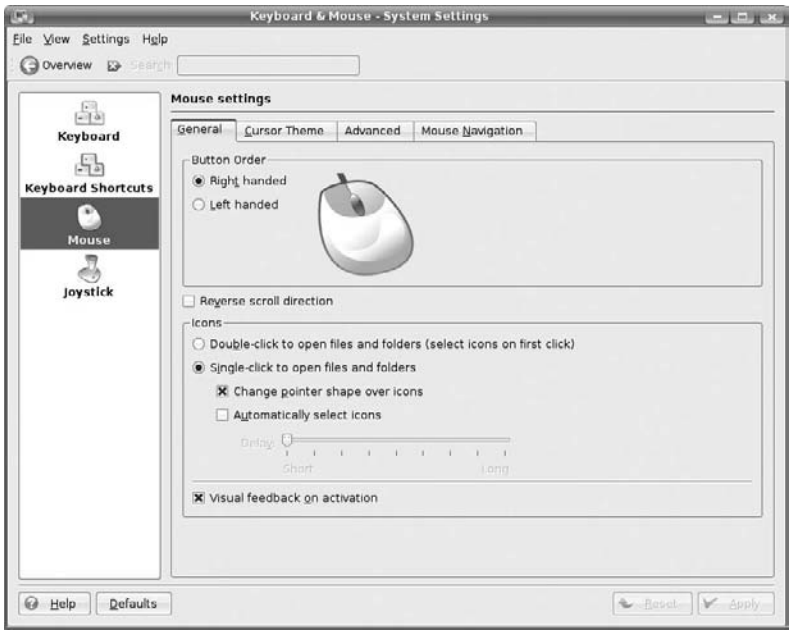
One of the most common customizations that people make is to change how responsive the mouse is, in one way or another. Available customizations include:

- Whether the mouse is configured for left-handed or right-handed users: The latter is the default, but may not suit you if you use your left hand to control the mouse.
- The type and size of the mouse pointer, which those of us who wear glasses often change, especially on laptop systems.
- The speed at which the cursor moves and the amount of screen real estate the mouse covers in relation to your moving the mouse.
- Whether a single-click or a double-click opens files and folders on the desktop or in a file manager.
- The period of time that can elapse between two mouse clicks to consider them as a double-click. If you have difficulty double-clicking but don't want to use a single click, you can increase the double-click time period.

KDE 3 makes it easy for you to configure these settings and more. To configure your mouse and keyboard settings, click the K Menu, select the System Settings application, and click the Keyboard & Mouse settings icon (located in the Computer Administration section). Click the Mouse icon in the left pane to display the configuration dialog shown in Figure 7-20.

FIGURE 7-20

The Keyboard & Mouse Settings dialog



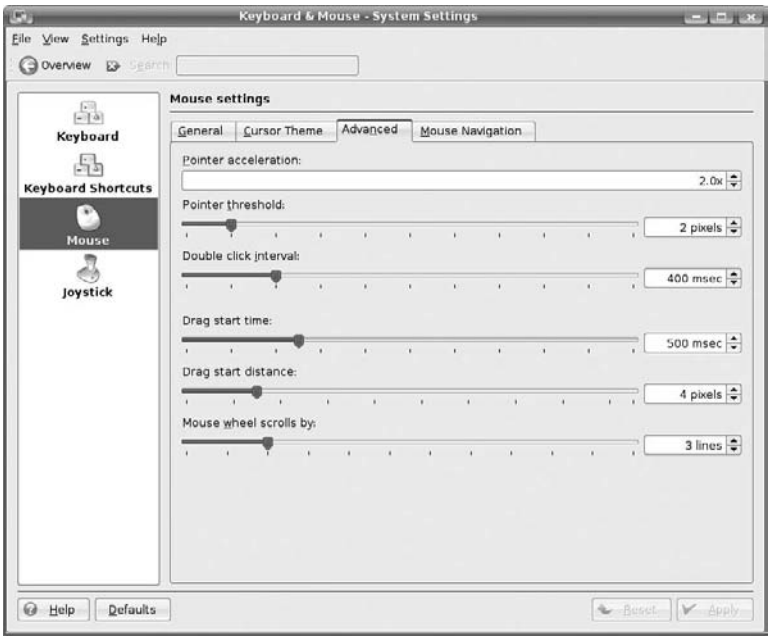
This settings dialog provides four tabs that enable you to configure different aspects of the mouse: the behavior of the buttons and how to select and execute programs via icons; what the mouse pointer/cursor looks like; advanced configuration options (where you can figure the speed and sensitivity of the mouse); and a mouse navigation tab where you can configure the speed and sensitivity of the mouse when you are controlling it using the arrow keys on a numeric keypad. The General tab is shown in Figure 7-20. To switch between different tabs, simply select the appropriate tab in the Keyboard & Mouse System Settings application.

TIP **One of the most frustrating things about KDE for new users is that, by default, KDE is usually configured so that a single click on a file or folder opens that file or folder rather than simply selecting it, as is done in almost every other graphical user interface I've ever seen. To disable this and return to the more standard "one click to select, double-click to open," select "Double-click to open files and folders" in the Icons section of the General Mouse Settings tab.**

The remainder of the configuration items in all of these tabs is pretty straightforward, so I won't explain each of them. However, the tab that does require some explanation is the Advanced tab, where you can configure the motion and sensitivity of the mouse. This tab is shown in Figure 7-21.

FIGURE 7-21

The Keyboard & Mouse Settings Advanced tab dialog



The entries on this tab have the following meanings:

- **Pointer acceleration:** This option controls the relationship between how quickly the cursor moves on the screen in response to your moving the mouse, and is set as a multiplier. The default value is that the mouse pointer moves twice as quickly on the screen as the mouse does. Being able to change this ratio guarantees that you don't need a mouse pad the size of your computer monitor—acceleration enables the cursor to cover larger distances as you move the mouse continuously. Slower values mean that cursor movement is closer to the actual distance that you're moving the mouse; faster values mean that the cursor moves proportionately farther distances as you continue to move the mouse—in other words, the cursor starts out moving slowly and accelerates as you continue to move the mouse. To change your settings, click the up or down arrows beside the current value to increase or decrease the acceleration factor value that you want to use.
- **Pointer threshold:** Related to pointer acceleration, this option lets you specify the number of pixels that the mouse must move in order to activate the acceleration multiplier. This value enables you to strike the right balance between making small, detailed mouse movements and being able to quickly move the mouse to other portions of the screen. If you are doing graphics or other onscreen work that requires tiny, detailed mouse movements, or if you are having problems controlling the mouse in general, you may want to increase this value.
- **Double-click interval:** Sets the maximum interval between two mouse clicks that should be interpreted as a double-click. If you are having problems double-clicking, you may want to increase this value.

- **Drag start time:** Sets the amount of time in which a single click can be interpreted as initiating a drag/move operation. Moving the cursor more than the distance specified in the Drag Start Distance parameter will initiate a drag operation.
- **Drag start distance:** Specifies the distance that you must move the mouse while an object is selected in order to begin dragging that object. You must move the cursor the specified distance within the time period specified by Drag Start Time in order for the drag operation to commence.
- **Mouse wheel scrolls by:** Only relevant if you are using a mouse with a scroll wheel, this option enables you to set the number of lines that an image should scroll for each click as you move the mouse wheel.

To make the changes to your mouse settings take effect, click Apply. (This button can be clicked only if you have actually made changes; otherwise, it is grayed out.) You can then experiment with the new settings and continue to change them until the mouse is as responsive (or unresponsive) as you would like.

To return to the main System Settings dialog, click the left arrow labeled Overview in the upper-left corner of the Keyboard & Mouse configuration screen.

Configuring Screen Resolution

KDE 3-based Kubuntu systems attempt to determine the size of your monitor, the maximum resolution that it can display, and the capabilities of your video card as part of the installation process. Unfortunately, this isn't always easy to do. Kubuntu therefore makes its best guess and provides a reasonable set of defaults, but also provides an easy-to-use System Settings application that enables you to manually configure these. Display resolution is one of those things that is hard to visualize in advance—it really makes sense only when you can see what different resolutions actually look like. Consequently, even if you are using a 30-inch monitor and the sexiest video card in the universe, you may find that your Kubuntu Linux system insists on displaying everything at a resolution of 1024×768, where each of the icons on your screen is approximately the size of a grapefruit. This is probably not what you had in mind, and is certainly not the best use of your hardware or screen real estate.

This section explains how to get the most out of your video hardware by selecting a higher screen resolution. Technically, the phrase "a higher screen resolution" means that your screen displays more pixels, but the visual effect is that things on the screen are both smaller and sharper. The folks in the computer industry, especially those who work at companies that produce graphics cards, generally refer to this as "the right thing."

Screen resolution in KDE is controlled by the K Menu ➪ System Settings ➪ Monitor & Display configuration application, which is shown in Figure 7-22.

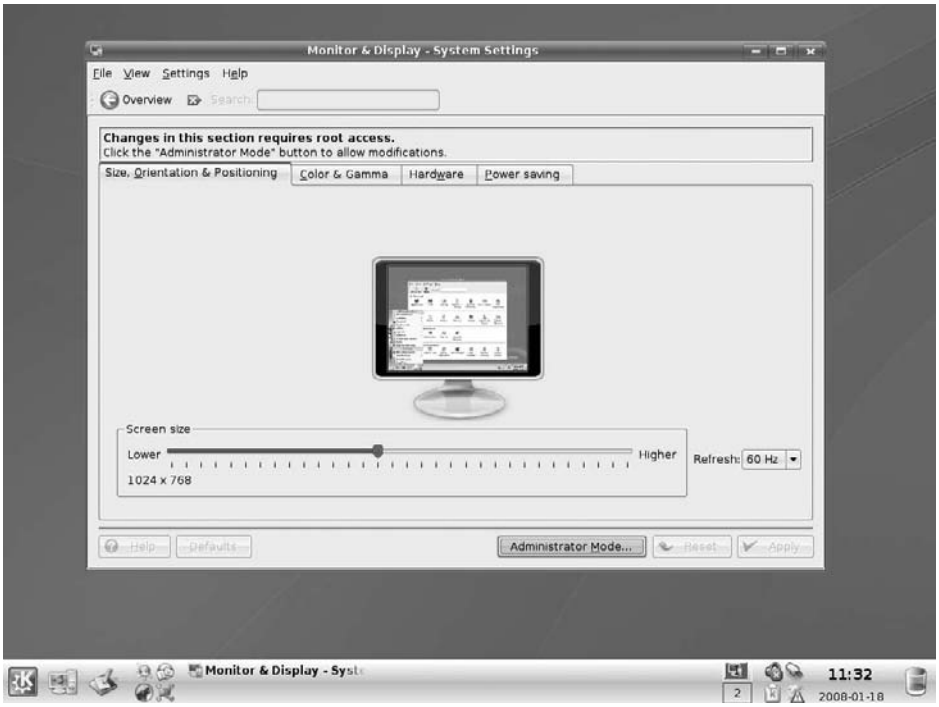
This application provides four tabs that logically organize different monitor-related configuration settings. The Size, Orientation & Positioning tab is displayed when you first start the application, as shown in Figure 7-22. The Screen size section at the bottom of the screen shows the current screen size, and enables you to select other screen sizes from those that are supported by your video card and monitor using a slider. The screen image above the Screen size slider shows a preview of what the screen will look like at the currently selected screen size.

You can also modify the Refresh rate using the control located to the right of the Screen size slider. This control enables you to select between different refresh rates (the rate at which the underlying X11 software redraws your screen) in order to reduce flicker and provide a more stable-looking screen. You can use the

up/down arrows to select a new refresh rate from those that are supported by your graphics hardware, or use your keyboard to enter a specific value. Higher numbers provide a more solid-looking display.

FIGURE 7-22

The Monitor & Display configuration application



To try any of these settings, click the Administrator Mode button and enter your password. Your system will redisplay the screen using the new screen size and refresh rate, and displays a dialog that asks whether you want to keep the new settings or revert to your previous ones. If you like the new settings, keep them—the new values will be automatically saved to the file `.kde/share/config/displayconfigrc` (in your home directory), and will be used any time that you subsequently log in on the system that you are currently using.

NOTE **Because it's possible to select a screen size that may not actually work correctly, your screen will automatically revert to its former resolution in 15 seconds. No need to panic if your screen goes haywire for a little while!**

The other tabs in the Monitor & Display configuration application enable you to configure the following:

- **Color & Gamma:** Enables you to use various sliders to configure how your system displays colors.
- **Hardware:** Enables you to provide detailed information about your monitor and video card, in case they were not detected correctly or they have additional capabilities that are not listed in the default information associated with them.

- **Power saving:** Enables you to activate power saving in your monitor if it is EnergyStar-compliant, and to set the period of inactivity after which your system will put the monitor into a low-power state.

Configuring screen resolution is probably the most common type of configuration that anyone does on any Linux system because it enables you to configure the system to match your aesthetic desires and suit the capabilities of your eyesight.

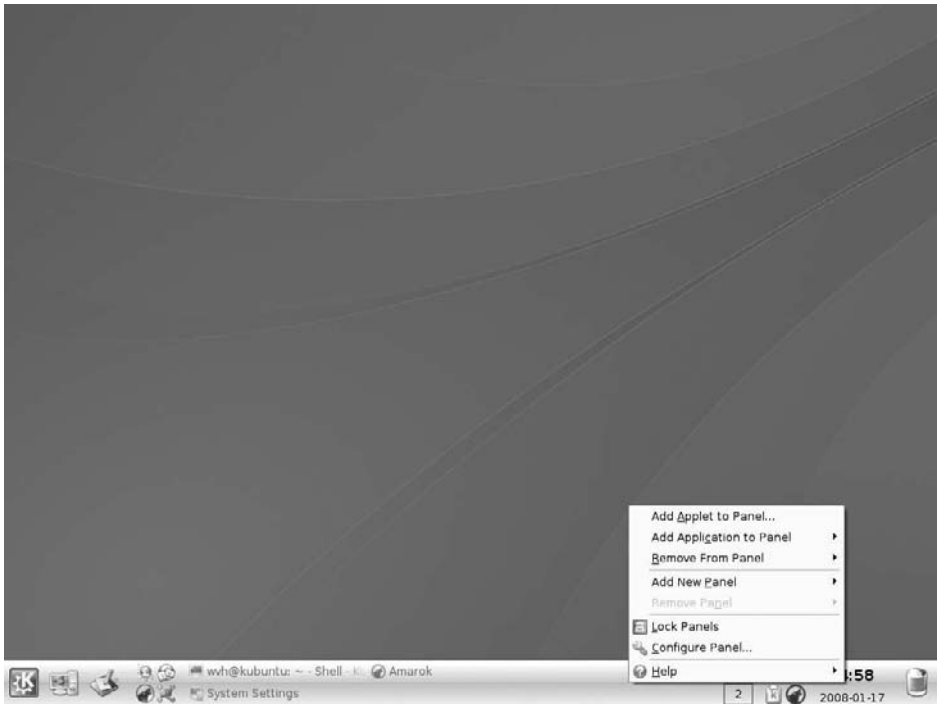
Customizing Panels

As discussed earlier in this chapter, a panel is a region on your desktop that provides access to various ways of interacting with the desktop and launching different types of applications. The KDE panel application (often known as the "Kicker") can display panels along any edge of the screen, supports multiple panels, and provides five different types of panels. The default KDE 3-based Kubuntu configuration displays a single panel at the bottom of the screen.

Right-clicking on any existing panel displays that panel's context-sensitive menu, as shown in Figure 7-23, which shows the context-sensitive menu for the default panel.

FIGURE 7-23

The KDE panel's context-sensitive menu



The context-sensitive menu for the KDE kicker panel displays the following configuration-related entries:

- **Add Applet to Panel:** Enables you to add a new applet, button, or other supported object to a KDE panel. See "Customizing Panel Contents," later in this section, for details on adding objects to an existing panel.
- **Add Application to Panel:** Enables you to add the launcher for a new application to a KDE panel. See "Customizing Panel Contents," later in this chapter, for details on adding objects to an existing panel.
- **Remove From Panel:** Enables you to select a currently installed panel applet and remove it from the panel.
- **Add New Panel:** Enables you to add a new panel to your KDE desktop. See "Adding Other Panels," later in this section, for details about the types of panels that are supported by KDE and how to add them to your KDE desktop.
- **Lock Panels:** Prevents panels from being modified by reducing the KDE panel's context-sensitive menu to contain only the "Unlock Panels" and Help entries. Selecting Unlock Panels restores the standard context-sensitive menu that is discussed in this section.
- **Configure Panel:** Displays a panel properties dialog that enables you to set certain properties for the currently selected panel. See "Customizing Panel Properties," later in this chapter, for details on the types of things that you can customize for an existing panel.
- **Help:** Displays the section of the online KDE documentation that discusses the KDE Kicker panel application.

The next three sections explain how to modify the properties of the default panel on your screen; how to add other applets, buttons, and controls to that panel; and how to add other panels to your screen.

Configuring the Panel

Selecting the Configure Panel entry on any standard KDE panel's context-sensitive menu displays the dialog shown in Figure 7-24.

This dialog provides access to five different configuration screens. After changing any of the values on any of these screens, click Apply to have them take effect so that you can examine their impact on your system. Clicking OK makes your change take effect but also closes the Panel Configuration application, which is probably not what you want to do if you are simply experimenting.

Figure 7-24 shows the primary configuration screen, the Arrangement screen, which is initially displayed when you select Configure Panel from the Kicker's context-sensitive menu. This screen enables you to specify the following:

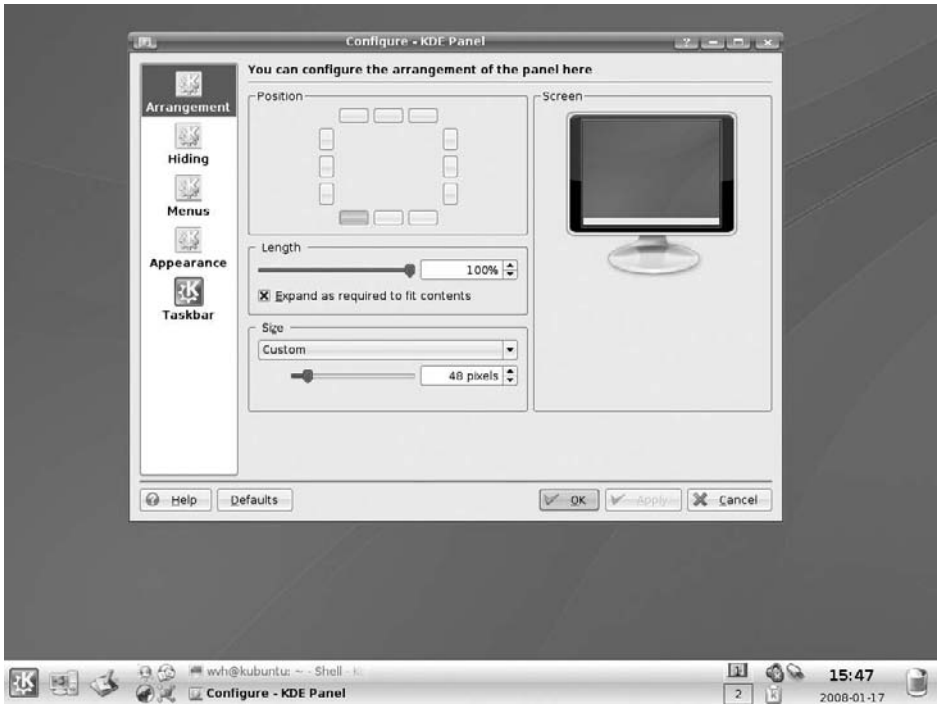
- **Position:** Select the default position of the panel by selecting any of the line segments around the rectangular outline that is provided. For panels that take up 100 percent of the side of the screen, there is no visible difference between selecting the left, center, or right of any edge of the screen—only the selected edge of the screen is visibly different. The screen image to the right of the outline shows a preview of what the screen will look like with the panel in the current position.
- **Length:** Enables you to specify the length of the panel as a percentage of the screen edge. To set this value, either click and drag the slider, use the up/down arrows to select a new percentage, or click on the current percentage and enter a new one using your keyboard. Setting the panel length to less than 100 percent will cause a visible difference between whether the panel is set to the left, center, or right of any edge of the screen. You can select the Expand as required to fit

contents checkbox if you want to initially use a smaller panel length, but have the panel grow, when necessary, if more space is required to display the current applets, tasks, and so on.

- **Size:** Enables you to set the height of the panel to any of four pre-selected heights (Tiny, Small, Normal, Large), or to specify a custom height using a slider, a numeric selector, or your keyboard.

FIGURE 7-24

The Kicker's panel configuration application



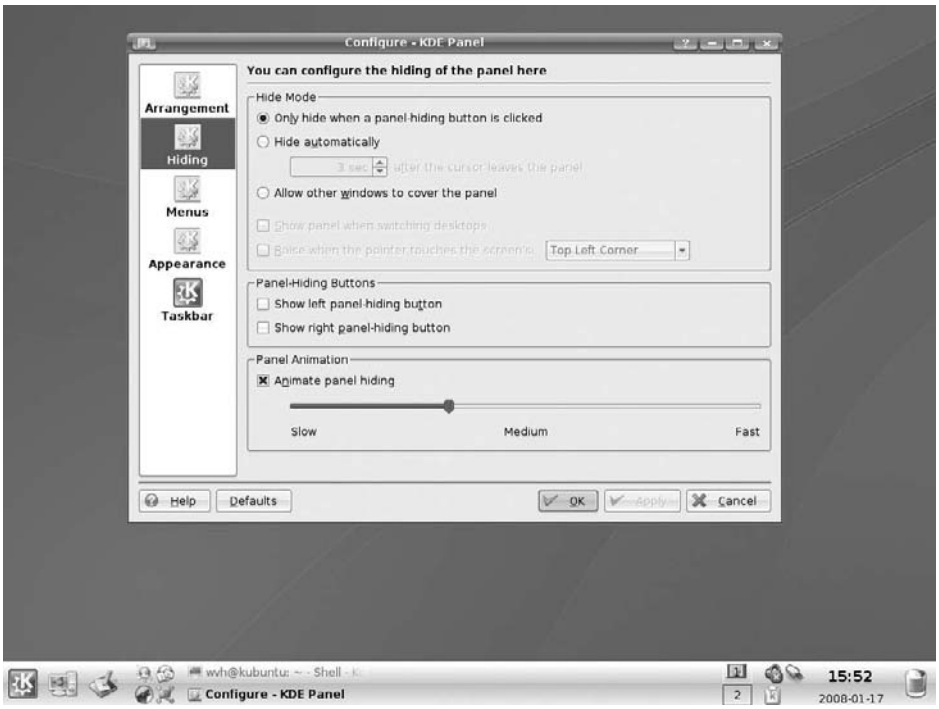
Remember that you must click the Apply button before any of your changes actually take effect on your screen.

Figure 7-25 shows the Hiding configuration screen, which is displayed if you click the Hiding icon in the left pane of Panel configuration application. These settings enable you to configure if (and when) the panel is hidden. Many people prefer to hide the panel when it is not actually in use, in order to maximize the amount of screen real estate that is available for viewing applications. (This is also a popular feature on Microsoft Windows and Mac OS X systems with regard to the Taskbar and the Dock, respectively.)

By default, the standard KDE panel is configured so that it is hidden only when you explicitly click a button to hide it. The buttons that you use to do this are not shown by default, which means that the panel cannot be hidden when its default configuration is used.

FIGURE 7-25

The Panel configuration application's Hiding screen



Automatically Hiding the Panel

To configure the panel so that it is automatically hidden:

1. Select the Hide Automatically radio button.
2. Use the up/down arrows to select the delay between the time that the mouse leaves the panel and when the panel is automatically hidden.
3. Click Apply to make these settings take effect.

If the panel is still configured to be displayed along the bottom edge of the screen, you will probably also want to select the "Raise when the pointer touches …" checkbox, and select Bottom Edge from the drop-down list. This will automatically raise the panel when you move the cursor to the bottom of the screen. You may also want to select the "Show panel when switching desktops" checkbox to ensure that the panel is initially displayed if you switch to another desktop using the Pager applet. Remember that you must click the Apply button before any of these additional changes actually take effect on your screen.

Figure 7-26 shows the Menus configuration screen, which is displayed if you select the Menus icon in the pane at the left of the Panel configuration application. As mentioned earlier in this chapter in the section entitled "Customizing KDE Menus," this configuration screen enables to you customize the contents of the K Menu by adding new, but pre-defined, items to the K Menu and configuring how those new items are displayed on the menu. The items that you can add are primarily new locations, but you can also activate two new sections of the K Menu using this configuration screen.

FIGURE 7-26

The Panel configuration application's Menus screen



The top (K Menu) portion of this screen enables you to add the following menu entries to the Actions portion of the K Menu:

- **Bookmarks:** Displays a submenu that provides easy access to any bookmarks that you have set in Konqueror.
- **Find:** Provides access to the KFind application (for desktop searches) and to Konqueror (for Web searches).
- **Kate Session Menu:** Makes it easy to start, resume, or restart the Kate (KDE Advanced Text Editor) editor.

- **Konqueror Profiles:** Simplifies starting Konqueror with a pre-defined profile. (See the section of this chapter entitled "Using Konqueror" for more information about Konqueror and profiles.)
- **Network Folders:** Enables you to launch Konqueror to display any of the network folders that you have configured.
- **Print System:** Provides easy access to applications for configuring and managing printing on your system, including the print server itself.
- **Quick Browser:** Provides quick links to various locations on your computer system—your home folder, the root of your system's filesystem—and links to configuration items in the `/etc` directory. While navigating any part of these hierarchies, you can open any directory location in the file manager, start a shell with the selected location as its working directory, or click on a specific file to open it using the application with which it is currently associated. The number of items displayed for each location in this menu is controlled by the value shown in the Quick Browser Menus section of this dialog.
- **Recent Documents:** Lists the most recently accessed documents on your system and enables you to automatically reopen them by re-launching the application that you were using to edit them.
- **Settings:** Enables you to start configuration applets used by the old KControl configuration application. KControl has been replaced by the System Settings application in later versions of KDE 3, so this menu option is fairly useless unless you explicitly install KControl and friends.
- **System Menu:** Adds the System Menu (available as a separate panel icon by default) to the K Menu. Adding this menu enables you to reclaim some panel space by removing the System Menu icon while still being able to take advantage of the convenient access to popular system locations that it provides.
- **Terminal Sessions:** Provides menu entries for launching various command-line applications: Shell, Linux Console, Python Interpreter, Midnight Commander, Root Shell, Screen Session, or new Konsole sessions with various pre-configured settings.

As always, you must click the Apply button before any of your changes will actually take effect on your screen.

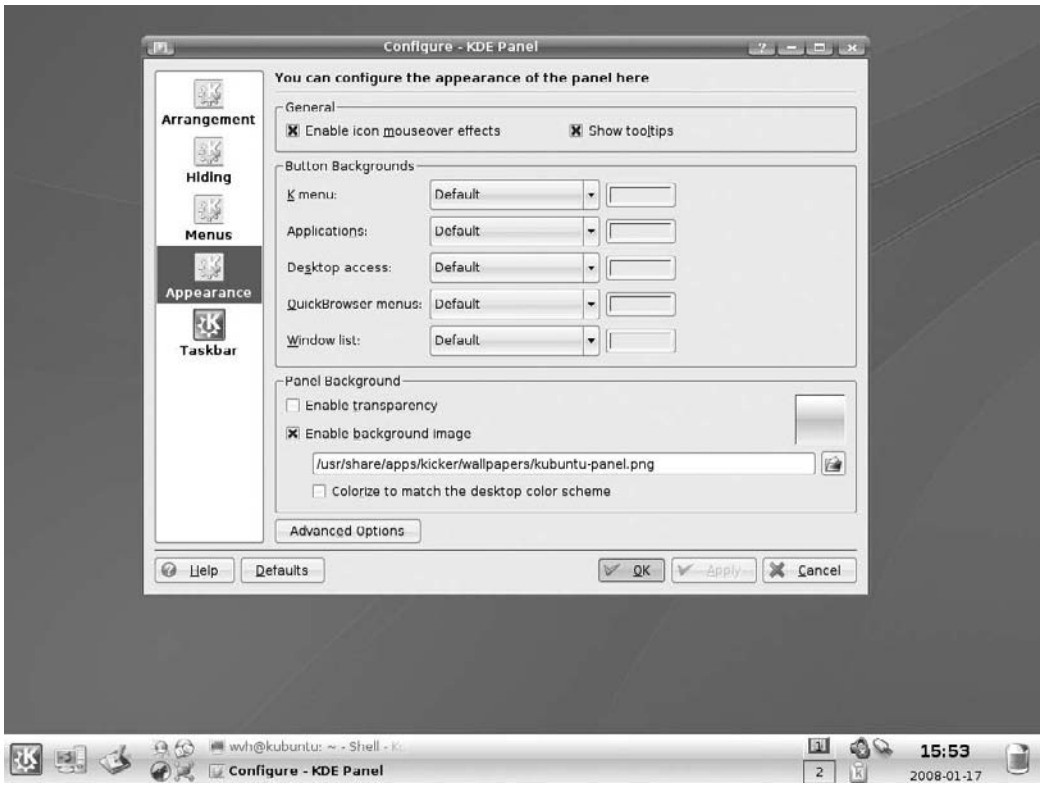
If the Quick browser menu is selected in the K Menu area at the top of this dialog, the QuickBrowser Menus area enables you to configure whether the Quick browser menus will show files whose names begin with a period (often referred to as "hidden files" because they don't show up in a normal directory listing) as well as the number of items that will be listed on any of the Quick browser menus. Some directories, such as the System Configuration directory (`/etc`), contain a large number of files, so you may find it useful to limit the number of files on each menu so that it fits on your screen.

The QuickStart Menu Items area enables you to configure the number of items that are shown on the Quick Start menu, which displays a list of the applications that you have most recently executed. Interestingly, this section of the K Menu is always active (and therefore isn't listed in the K Menu area at the top of this screen), but because the default number of applications that it shows is 0, it isn't visible on the K Menu. Setting the Maximum number of entries to any non-zero value causes this section of the K Menu to become visible. By default, this section lists the applications that are most frequently used. You can select the Show the applications most recently used radio button to configure this menu to show the applications that you have most recently used, rather than those that you use most often.

Selecting the Appearance icon in the Panel configuration application's left pane displays the screen shown in Figure 7-27. This screen enables you to configure various aspects of the appearance and visual behavior of the panel.

FIGURE 7-27

The Panel configuration application's Appearance screen



The Appearance panel enables you to configure the following settings:

- **Enable icon mouseover effects:** Configures whether moving the mouse over any portion of a panel applet causes a relatively large explanatory image to display. This image identifies the name of the panel applet and displays summary information about its purpose or status. I personally find this irritating, so I typically disable this because the tooltips (see the next item) generally provide enough information for me.
- **Show tooltips:** Configures whether moving the mouse over any portion of a panel applet causes a tooltip to display. Tooltips are yellow rectangular boxes that display summary information about the system object that is associated with an icon (a panel applet, in this case). Tooltips are generally useful without being overly intrusive (which I find the previous item, the mouseover effects, to be). Your mileage may vary.
- **Button backgrounds:** Configures the background used when displaying various types of icons in the panel. This can be useful if you want to visually distinguish a frequently used applet in order to locate it more easily. To change the icon used for any icon in the list, click the drop-down

list of available background shapes and colors, select a new one, and click Apply to have your changes take effect.

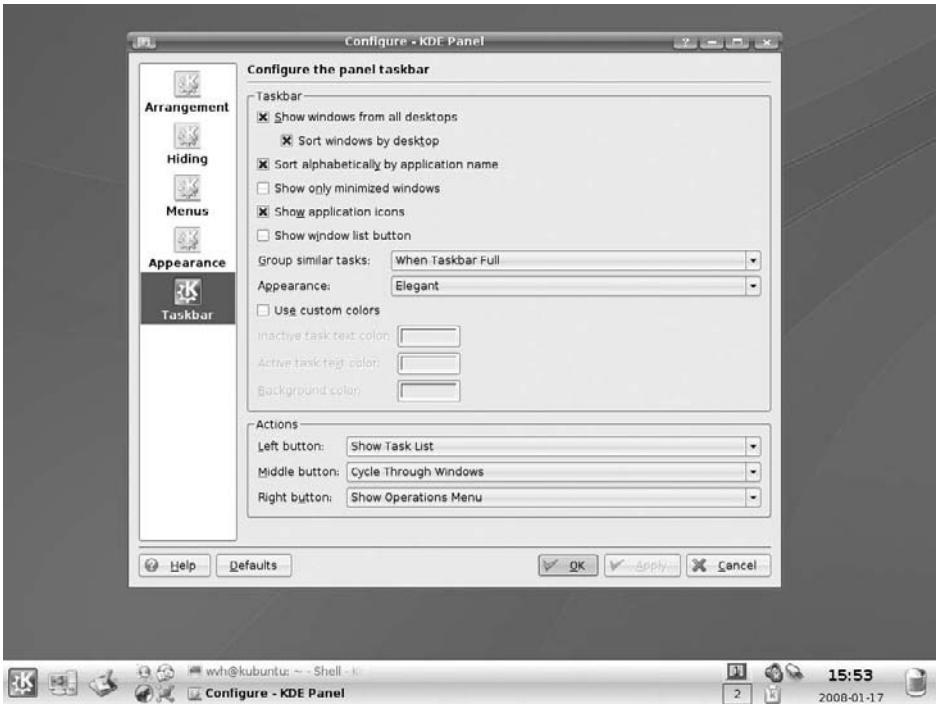
- **Panel backgrounds:** Configures whether the panel is transparent and, if not, the image that is used as its background. Making the panel transparent disables the use of a background image, and can be useful if you don't want to hide the panel but find its appearance to be too intrusive. Whether you want to change these values depends on your particular sense of aesthetics.

As always, you must click the Apply button before any of your changes will actually take effect on your screen.

Clicking the Taskbar icon in the left Panel Configuration pane displays the screen shown in Figure 7-28. This screen enables you to configure various aspects of the appearance and behavior of the KDE Panel's Taskbar applet, which is the applet that displays icons for each of the applications that are currently running on the desktop.

FIGURE 7-28

The Panel configuration application's Taskbar screen



After the K Menu icon, the Taskbar applet is probably the most useful of the KDE Panel applets because it provides easy access to the applications that you are running. Clicking on any application's entry in

the Taskbar minimizes or maximizes the associated application, based on its current state. The Panel Configuration application's Taskbar screen enables you to configure the following aspects of the Taskbar:

- **Show windows from all desktops:** Determines whether the Taskbar contains entries for all applications that are running on any of your desktops. If you want the Taskbar to show only applications that are running on the current desktop, deselect this option. This option is enabled by default on KDE 3-based Kubuntu systems, as is the Sort windows by desktop option, which lists applications from desktop 1 first, desktop 2 next, and so on.
- **Sort alphabetically by desktop:** Sorts the applications in the Taskbar alphabetically (within each desktop if the Sort windows by desktop option is selected).
- **Show only minimized windows:** Configures the Taskbar so that it contains entries only for applications that have been minimized. Applications that are currently displayed on a desktop are not listed in the Taskbar. This option is disabled by default, which seems like the right thing to me because I find it useful to be able to left-click Taskbar entries to minimize/maximize applications, and you can't left-click on anything that isn't listed.
- **Show application icons:** Determines whether the icon for an application is displayed to the left of the name of that application when it is shown in the Taskbar. De-selecting this option reduces the size of Taskbar entries, which many people find to be useful (myself included).
- **Show window list button:** Displays an up-arrow icon that you can click on to display a menu that lists all of the applications that are listed in the Taskbar, sorted into regions based on the desktop on which they are running.
- **Group similar tasks:** Enables you to configure whether each instance of an application gets its own icon in the Taskbar, or whether all instances of a single application are grouped together under a single icon that you must select to minimize/maximize a specific application. The default value is When Taskbar Full, which activates this feature only when you have enough applications running on your system that the Taskbar is full of application icons. Other choices are Never and Always. I generally leave this at its default value because I rarely fill up the Taskbar and this makes sorting Taskbar entries by desktop more complex, but this option ensures that the typical KDE flexibility is there if you run zillions of applications simultaneously.
- **Appearance:** Enables you to configure how applications are shown in the Taskbar.
- **Use custom colors:** Enables you to assign different colors to inactive and active application entries in the Taskbar.
- **Actions:** Enables you to select the actions associated with each mouse button when you click in the Taskbar.

As always, when experimenting with any of the values on this screen, you can click Apply to see their impact. When you are done configuring panel applications, you can click OK to save all of the changes that you have made and close the Panel Configuration application.

Customizing Panel Contents

Configuring the KDE panel and configuring the contents of that panel go hand in hand. The Kicker's context-sensitive menu makes it easy for you to add applets to or remove applets from the panel by selecting appropriate menu options and using the associated dialogs and submenus.

To add applets to the KDE panel, right-click on the panel to display its context-sensitive menu. Next, select Add Applet to Panel to display the Add to Panel dialog, as shown in Figure 7-29.

FIGURE 7-29

The Add to Panel dialog



This dialog enables you to add application launchers and common applets to any Kicker panel that is currently displayed on your screen. All of the currently available applets are displayed in this scrollable dialog. Applets that are currently located on the panel are lightly highlighted, as the Clock applet is in Figure 7-29. To add any of these to any panel that is currently visible on the screen, left-click on that applet to select it and click the Add to Panel button. (You can click this button only once you have selected an applet—it is grayed-out until then.) You can then move it as described later in this section.

On KDE 3-based Kubuntu systems, the applets that are available in this dialog by default are as follows:

- **Bookmarks Menu:** Provides access to any bookmarks that you have defined in the Konqueror object browser. See the section "Using Konqueror" later in this chapter for more information about these bookmarks.
- **Clock:** Displays an analog or digital clock. This applet is shown in the KDE panel on Kubuntu systems by default.
- **Desktop Preview & Pager:** Enables you to manage and switch between multiple workspaces and provides a miniature display of each active workspace. Workspaces are essentially separate virtual screens that are provided by the KDE desktop. You can switch between virtual desktops by clicking on them in the applet, start applications on different workspaces, move applications between workspaces by dragging and dropping in the preview, and so on. Multiple workspaces provide a convenient way of running different types of applications on different virtual screens without them being visible until you actually switch to the workspace where they are displayed. This is useful, for example, when you're at work and want to play a game—you can start the game on a different workspace from the one where you're actually working, and switch to it whenever no one is looking over your shoulder. This applet is shown in the KDE panel on Kubuntu systems by default.
- **Find:** Enables you to easily search files or Web locations for specified terms or phrases.
- **Kate Session Menu:** Enables you to open the Kate (KDE Advanced Text Editor) with pre-configured files or new ones.
- **Keyboard Status Applet:** Displays the state of various keyboard modifier keys (Shift, Ctrl, Caps Lock, and so on) and the mouse buttons.
- **Klipper:** Provides a centralized clipboard for all KDE applications. This applet is active in KDE on Kubuntu systems by default; an icon that provides access to Klipper is always visible in the Notification Area/System Tray.
- **K Menu:** Displays the default menus and menu hierarchies associated with KDE on Kubuntu systems. This applet is shown in the KDE panel on Kubuntu systems by default.

- **Konqueror Profiles:** Provides easy access to Konqueror, enabling you to start Konqueror for file management or Web browsing, or to use a profile that you have configured.
- **Lock/Logout Buttons:** Displays buttons that enable you to lock the screen or log out with a single mouse click.
- **Network Folders:** Displays a menu that provides easy access to default network locations (Samba Shares, Network Services) and to specific network shares that you or your system administrator has defined
- **Non-KDE Application Launcher:** Simplifies the creation of a button that will execute a specified application. As discussed later in this section, the Kicker's context-sensitive menu makes it easy for you to add an icon which launches an application that is already on the K Menu. This applet provides a simpler, but more powerful, interface for creating a launcher that will start any command with a specific set of options in response to selecting an icon.
- **Print System:** Displays a menu that gives you easy access to applications for configuring and managing printing on your system, including the print server itself.
- **Public File Server:** Enables you to start a small, lightweight Web browser and use it to share files and folders from a specified directory.
- **Quick File Browser:** Provides quick links to files and subdirectories in a specific directory on your system. When displayed in the panel, this applet is frequently used to provide quick, graphical access to a user's home directory.
- **Quick Launcher:** Makes it easy to launch your favorite applications by adding an icon for each of them to this applet's display area in the panel. This applet is shown in the KDE panel on KDE 3-based Kubuntu systems by default.
- **Recent Documents:** Provides easy access to a list of the documents that you have worked with most recently, and enables you to easily reopen those documents in their associated applications.
- **Runaway Process Catcher:** Starts a process and associated display area that notifies you when an application is taking too much CPU time. This applet makes it easy for you to notice and deal with applications that might otherwise affect the overall performance of your system.
- **Run Command:** Provides a small dialog in which you can enter a specific command that you want to execute. If you are not a daily user of Linux command-line tools, running this applet may be more convenient for you than always running a shell "just in case."
- **Settings:** Provides easy access to the KDE Control Center, which is a KDE configuration application that has been replaced by the System Settings application discussed earlier in this chapter. This entry is provided to simplify access to legacy settings applications.
- **Show Desktop:** Provides quick access to the KDE desktop by minimizing existing applications to expose the desktop. This applet is shown in the KDE panel on Kubuntu systems by default.
- **Skim Panel Applet:** Provides an applet that is required if you are running applications that are SCIM-aware (Smart Common Input Method) and want to be able to dock those applications on the panel. SCIM provides a flexible mechanism for handling different input sources, character sets, and so on.
- **Sound Mixer:** Provides easy access to a specified sound mixer, which is an application that provides volume and level settings for various audio cards and audio subsystems. Popular graphical sound mixers on Kubuntu are alsamixer and kmix. This applet is active in KDE on Kubuntu systems by default; an icon that provides access to your sound mixer is always visible in the Notification Area/System Tray.

- **Storage Media:** Provides easy access to removable or networked storage media that are currently connected to your system.
- **Strigi:** Provides an icon and small text area for using the Strigi application to perform a desktop search on your system, much like Spotlight on Mac OS X or Microsoft Windows desktop search application. See the section "Using the Strigi Desktop Search Application," later in this chapter, for more information about Strigi.
- **System Menu:** Provides an icon and associated menu that gives easy access to specified locations on your Kubuntu system. This applet is shown in the KDE panel on KDE 3-based Kubuntu systems by default.
- **System Tray:** Provides a display area in which applications and applets can display icons that represent system messages (such as the arrival of e-mail or low-power conditions) or provide access to system resources such as the clipboard used for cutting and pasting data between applications. The system tray concept should be very familiar to users of other graphical interfaces, such as Microsoft Windows. One important application that shows messages here is the Adept Updater, which displays an icon in the notification area to let you know that system updates are available.
- **Taskbar:** This applet is shown in the KDE panel on Kubuntu systems by default.
- **Terminal Sessions:** Provides an icon that enables you to launch various command-line applications: Shell, Linux Console, Python Interpreter, Midnight Commander, Root Shell, Screen Session, or new Konsole sessions with various preconfigured settings.
- **Trash:** Provides the familiar "desktop trashcan" as a central location for access to and the intermediate storage of deleted files and directories. This applet is shown in the KDE panel on Kubuntu systems by default.
- **Window List Menu:** Provides an icon that enables you to launch various command-line applications: Shell, Linux Console, Python Interpreter, Midnight Commander, Root Shell, Screen Session, or new Konsole sessions with various preconfigured settings.

Many of these applets require some initial configuration before you can use them. This is usually done in one of two ways:

- **Via the applet handle:** Most applets provide an invisible handle to the left of the icon that provides access to an applet configuration and control menu. To see this handle, move the mouse just to the right of the applet's icon, which should expose a thin grey area with an up arrow at its top. Right-clicking on this handle should display a menu that contains a "Configure applet-name" entry.
- **By right-clicking on the applet icon:** Applets without a handle typically provide a configuration option on their context-sensitive menu.

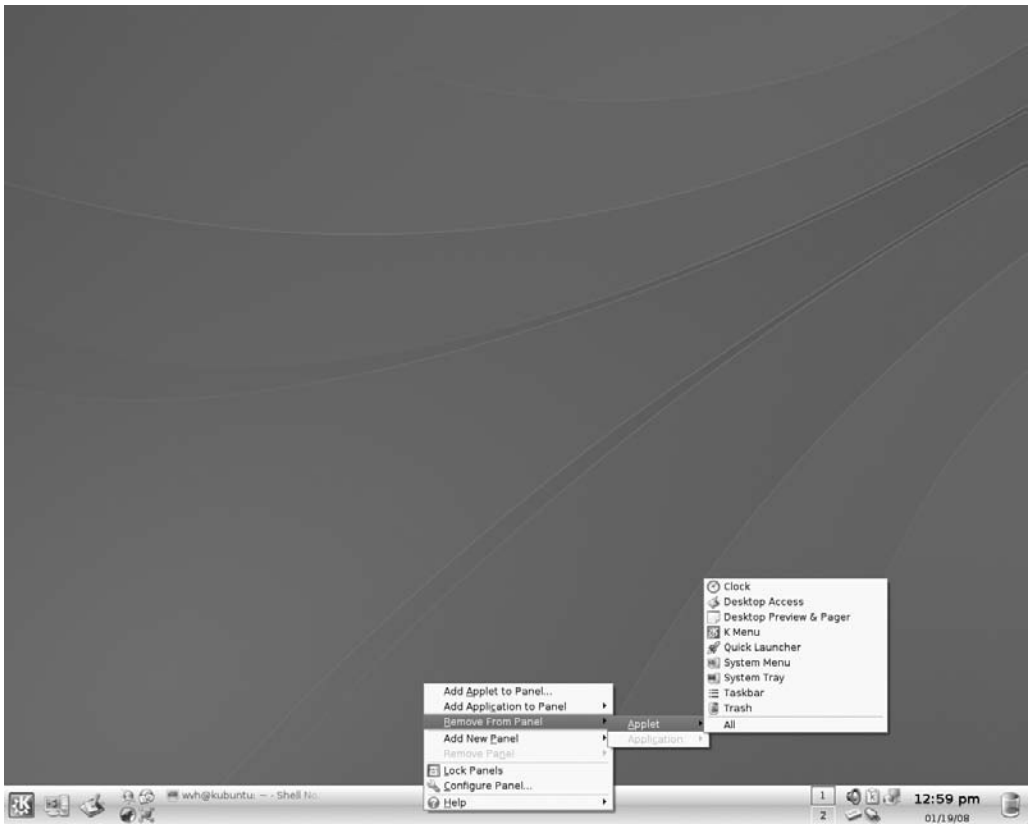
Moving applets is done in the same way as configuring them: A Move *applet-name* menu option is typically provided through the applet handle or through its context-sensitive menu.

Removing applets is done in one of three ways:

- By selecting "Remove *applet-name*" from the applet handle's context-sensitive menu
- By selecting "Remove *applet-name*" from the applet's context-sensitive menu
- By right-clicking anywhere on the panel, and selecting the Remove from panel ➪ Applet ➪ Applet-name, as shown in Figure 7-30.

FIGURE 7-30

Removing an applet from the panel



The KDE Kicker panel application also makes it easy for you to add icons that you can click on in order to start specific applications. You can do this in either of two ways:

- Add a launcher for any item that is listed on the K Menu by right-clicking on the panel, selecting the Add Application menu item, and navigating to the menu item that is associated with the application for which you want to add an icon.
- Add a launcher for any non-K-Menu item on your system by either adding and configuring the Non-KDE Application Launcher applet, or by selecting Add Application ➪ Non-KDE Application from the panel's context-sensitive menu.

Adding Other Panels

As mentioned previously, you can display more than one panel (and more than one type of panel) on the screen at the same time. All of the panels that are displayed on your screen are controlled by a single instance of the kicker application—the fact that multiple panels are displayed on your screen does not

mean that multiple copies of the kicker are running. You can add new panels of the following types to your KDE desktop:

- **Dock application bar:** A panel that enables users to dock applications that were designed for the WindowMaker window manager. Not that this comes up all that often nowadays …
- **External Taskbar:** A dedicated taskbar that, because it is larger, provides space for seeing and tracking more applications than the standard Taskbar that is embedded in the default panel. Using this Taskbar also frees up space for running other panel applets in the default panel.
- **Kasbar:** An alternate taskbar that was inspired by the Taskbar used by window managers such as WindowMaker and Black Box.
- **Panel:** Another instance of the standard KDE Panel.
- **Universal sidebar:** An alternate taskbar that was originally a part of the Konqueror browser. This can serve as a filesystem and home folder browser, a network location manager, a history manager for visited folders and sound files played via Amarok, a convenient location for storing desktop and system bookmarks and whatever else you have to store.

To add a new panel to your system, right-click on the initial KDE panel, select the Add New Panel submenu, and select the type of panel that you want to add from the menu that displays. To remove an existing panel, right-click on the initial KDE panel, select the Remove Panel submenu, and select the panel that you want to remove from the menu that displays.

By default, the Kicker panel displayed at the bottom of your screen contains the menu icons, applets, and other items discussed in the "KDE Desktop Overview" section earlier in this chapter. Not only can you configure how and where the panels are displayed on your screen and how many panels are displayed, but you can also extensively customize the contents of each. Kubuntu's default KDE installation comes with many other applets, buttons, and controls that you can install in the panel on your desktop.

Configuring a Screensaver

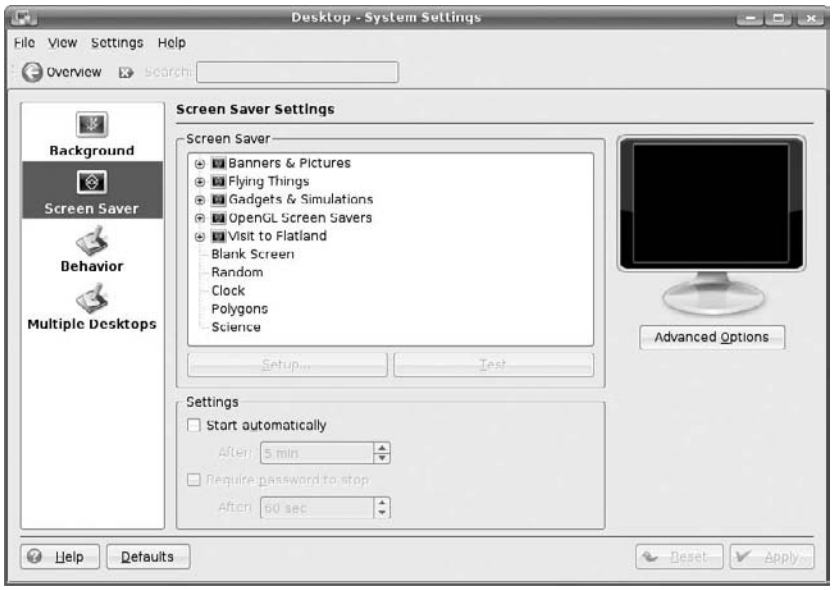
There was a time when computer monitors were quite sensitive to displaying the same characters for long periods of time. Monitors that showed the same applications in the same locations for long periods of time would end up damaging the phosphors on the inside of the monitor, which therefore displayed a ghostly outline of that application on the screen. (This was known as screen burn or video burn.) To eliminate, or at least, reduce this problem, many companies developed applications known as screensavers, which would start automatically after a computer system had been unused for a specified period of time, and would display random patterns or images on the screen, preventing the burn-in of static images.

Today's monitors are much more sophisticated than the monitors of old, and are therefore less susceptible to image burn-in. Today, screensavers are typically used because they look nice or for security purposes, so that unauthorized users can't see what is displayed on the screen of your computer while you're at lunch.

By default, your KDE 3-based Kubuntu system comes configured to blank out the screen after 10 minutes of keyboard and mouse inactivity. As an alternative to simply turning the screen off, you can select a specific screensaver for your system to use (hundreds are available) and configure the period of time after which a screensaver is activated. All of these are done through the K Menu ➪ System Settings ➪ Desktop configuration application. After starting this configuration application, select the Screen Saver icon from the list at left. The Screen Saver Settings screen displays, as shown in Figure 7-31.

FIGURE 7-31

Configuring your system's screensaver



Once this screen displays, you can explore all of the other screensavers that are available on your system by selecting them from the scrollable Screensaver pane. A preview of the appearance and behavior of the currently selected screensaver displays in the preview pane to the right of the screensaver list.

Once you've selected the screensaver that you want to use, click Setup to display a dialog that enables you to set some general characteristics of your system's use of that screensaver.

Once you've selected and configured your screensaver, click Test to see it in action. Clicking anywhere on the screen terminates this test run and returns you to the Screen Saver Settings screen.

Once you are happy with your choice of screensaver and have optionally fine-tuned it via the Setup dialog, click the Start automatically checkbox to enable the use of the screensaver. You can also enter a new period of time after which the screensaver should start. The default time period is 5 minutes. You can also check the Require password to stop checkbox if you want to require the current user to type their password in order to terminate the screensaver and redisplay the screen. If you enable this option, you can also specify the time period before which no password is required. This provides a grace period that enables you to notice that your screensaver has activated but terminate it by moving the mouse rather than having to enter your password each time that it starts.

Once you are finished configuring the screensaver, click Apply to save your changes permanently. Click the overview button in the heading of this screen to close the Screen Saver Settings screen and return to the top level of the System Settings application.

Configuring Power Management

Power management refers to mechanisms that enable your system to reduce or minimize power consumption in various cases, the most common of which are after a system has been idle (i.e., no keys have been pressed and no disk activity has occurred) for a specified period. The most common power management mechanisms do things like temporarily turning off power to your monitor, spinning down hard drives, suspending your system by saving its current memory and application state to RAM or unused portions of your hard drive, and so on. See the "Configuring Power Management" section of Chapter 5 for a discussion of general power management terminology.

The best power management utility for KDE 3-based systems is an applet known as KPowersave. It is not installed by default on KDE 3-based systems, but that can easily be corrected by using your favorite package management utility from those discussed in Chapter 23 to install it on your Kubuntu 8.04 or earlier system. (KDE 4-based systems support power management through applications called Guidance or PowerDevil, depending on the version of KDE 4 that you're using, and KPowersave seems to have been phased out.)

TIP **The System Settings ➪ Monitor & Display configuration tool also supports simple settings for managing power use on your monitor. This configuration tool provides a Power Savings tab that enables you to activate power saving on your monitor and specify the period of time after which your system will turn off power to your monitor. If this is the only aspect of power management that you're interested in, you can simply use this configuration tool, and needn't worry about installing or using KPowersave. However, you can do much more with KPowersave—whether you choose to do so is up to you.**

After installing KPowersave, you can start it by selecting the K Menu ➪ System ➪ kpowersave - Battery Monitor utility menu command. The name is slightly misleading because KPowersave doesn't simply monitor and manage power consumptions if you are running from battery power, but helps you monitor and manage power consumption and savings on any type of system. Portable, battery-powered systems are the traditional area in which power management is most critical, but if you leave your desktop and server systems running for long periods of time and your electric bill is anything like mine, any savings in power consumption is much appreciated.

After starting KPowersave, its icon (the plug at the end of a power cord) appears in the panel's System Tray. You can left-click this icon at any time to display a KPowersave dialog that shows the current power management scheme that KPowersave is using, the status of the HAL (Hardware Availability Layer) daemon on your system (which KPowersave uses to interact with your hardware), and information about the CPU(s) on your system and their speed, usage, and so on. While this is interesting, the most important aspect of using KPowersave is that you can cause specific actions to occur or configuring it to match the specific requirements of the system on which it is running. Right-clicking the KPowersave icon in the system tray displays a menu that provides the following options:

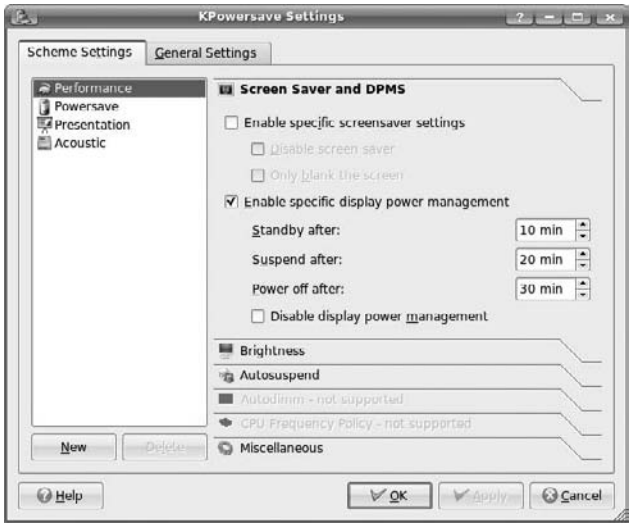
- **Configure KPowersave:** Displays KPowersave's Configuration dialog, which is discussed later in this section.
- **Configure Notifications:** Displays a dialog that enables you to configure the notifications that are associated with different power-related events such as switching power management schemes, suspending to disk or RAM, and so on.
- **Suspend to Disk:** Immediately begins suspending the system to disk, which completely powers off your system once the process completes.
- **Suspend to RAM:** Immediately begins suspending the system to RAM, which powers off the monitor and spins down your disks once the process completes.

- **Set Active Scheme:** Provides a quick mechanism for switching between different defaults of user-created power management schemes. Power management schemes are defined and configured on the Scheme Settings tab of the KPowersave Configuration dialog.
- **Help:** Displays the KPowersave manual in the Help document/help viewer.
- **Quit:** Terminates the KPowersave applet.

To configure KPowersave, right-click its icon in the System Tray and select the Configure KPowersave menu option to display a configuration dialog like the one in Figure 7-32.

FIGURE 7-32

Configuring power management in KPowersave



The KPowersave configuration dialog provides two tabs:

- **Scheme Settings:** A tab that lets you create and manage power management schemes, where a scheme is simply a collection of power management settings with a single name that make it easy for you to change multiple settings at the same time.
- **General Settings:** A tab that enables you to specify your default power management scheme when running off wall or battery power, define the actions that KPowersave should take when specific system events occur (such as pressing the power button or closing the lid of a laptop), whether to lock the screen when turning it off, and so on.

The Scheme Settings tab displays by default when you are configuring KPowersave. The list at the left of this dialog shows predefined power schemes. The power management scheme that you are currently using is selected by default. To switch to a different power management scheme, select its name in this list and click OK.

The portion of this dialog at the right displays groups of related settings that you can expand by clicking their names. The groups of power-related per-scheme settings provided by KPowersave are as follows:

- **Screen Saver and DPMS:** Settings for how (and if) you use a screensaver, and general Display Power Management Signaling (DPMS) settings such as standby, suspend, and poweroff.
- **Brightness:** Settings for adjusting the brightness of your monitor (not supported on all hardware).
- **Autosuspend:** Settings for enabling the suspension of the system to RAM or disk based on idle time.
- **Autodimm:** Settings for dimming your monitor based on idle time (not supported on all hardware).
- **CPU Frequency Policy:** Settings for changing the CPU frequency based on current usage and idle time (not supported on all hardware).
- **Miscellaneous:** Settings for enabling notification based on power management events.

The values that are shown in these groups as you expand them are those associated with the currently selected scheme from the list at left.

The General Settings tab enables you to configure power management settings that are independent of any specific power management scheme, and contains groups of related settings that you can expand by clicking their names. The groups of general settings provided by KPowersave are as follows:

- **Screen Saver and DPMS:** General settings for screensavers on the system, and general Display Power Management Signaling (DPMS) settings (not supported on all hardware).
- **Button Events:** Enables you to map the system Shutdown, Logout, Suspend to Disk, and Suspend to RAM events to specific buttons, switches, or events on a system, such as pressing the Power Off button, closing the lid of a laptop, pressing the Sleep button, or pressing the Suspend-to-Disk button.
- **Default Schemes:** Enables you to set the default power management schemes used when a system is running on AC or battery power.
- **Lock Screen:** Enables you to configure whether the system automatically locks the screen before suspending to disk or RAM, and the software mechanism used to do so.
- **Miscellaneous:** Enables you to configure power-related system notifications, set whether KPowersave starts automatically when you boot your system, and define applications that cannot be running if suspending a system to disk or RAM is requested. You must terminate these applications before you can suspend your system. The blacklists are provided to prevent conflicts between automatic suspend timeouts and long-running applications.

After modifying any values on the Scheme or General Settings tabs, click OK to save your changes and close the KPowersave configuration dialog, or click Apply to save your changes and continue configuring KPowersave.

See the "Configuring Power Management" section of Chapter 5 for suggestions on how to debug and (one hopes) correct any power management problems on your systems.

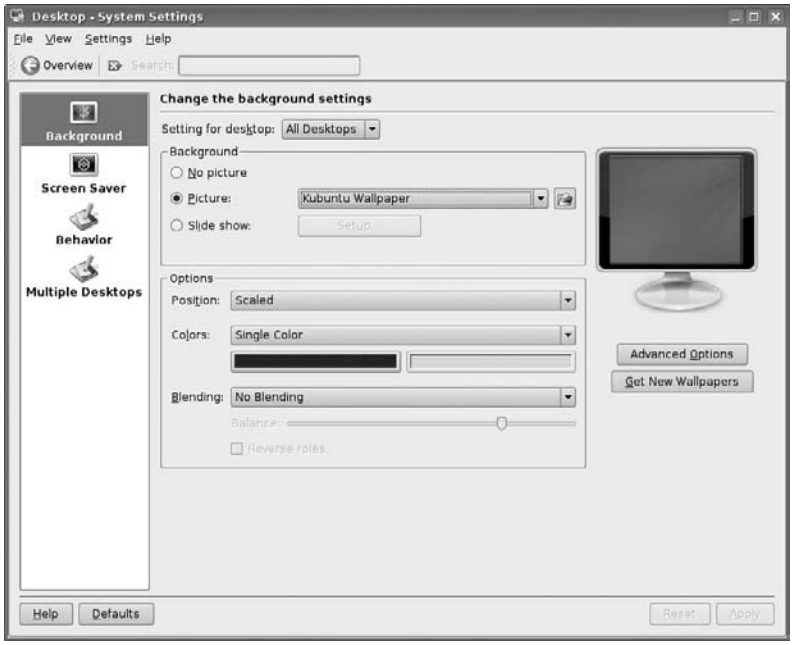
Changing Desktop Backgrounds

Changing the background of your desktop is one of the more common configuration tasks that anyone wants to do when personalizing a system. KDE's System Settings application makes it easy for you to change the background displayed on your system's desktops. The background is a solid color or image that is displayed on your screen and over which all other application windows are displayed. A background is frequently referred to as the "wallpaper" for your desktop.

To change your desktop background, start the K Menu ➪ System Settings application and select the Desktop icon in the Look & Feel section. The Background screen displays, as shown in Figure 7-33.

FIGURE 7-33

Configuring your desktop background



To customize your desktop, you can specify that no image is used (using a solid background color instead), select an existing graphics image from the default list that is provided in the Picture list, or click the Browse icon beside the list of available pictures to navigate to the location of a specific graphics file that you want to use as your background. (I tend to use vacation or other scenic photos, myself.) Sometimes it's necessary to remind yourself that there's more to life than just a computer screen—and what better way to do that than by displaying something from real life on your screen? I wish I were kidding.

Once you've selected an image, the Options section's Position value determines how that image is used as the background. Your options are as follows:

- **Centered:** The selected image is displayed in the center of your screen at its original resolution and size.
- **Tiled:** The selected image, at its original size, is placed in the upper-left corner of the screen and is repeated across and down your screen as many times as needed to fill the background.
- **Center Tiled:** Like Tiled, except that the selected image is placed in the center of your screen and then is replicated, up, down, and to either side as many times as needed to fill the background.
- **Centered Maxpect:** The selected image is placed at the center of your screen and is scaled to fit your desktop without changing the original aspect ratio of the image (as in a movie or television program shown in letterbox format). Any space that cannot be filled without distorting the image will be black.
- **Tiled Maxpect:** Like Centered Maxpect, except that any space that cannot be filled without distorting the image will be filled with copies of the original image.
- **Scaled:** The selected image is stretched to completely fill your screen, changing its aspect ratio if necessary.
- **Centered Autofit:** Like Centered, except that if the image is larger than your current desktop, it will be scaled down to best fit your screen without changing its aspect ratio.
- **Scale & Crop:** The selected image is magnified (without changing its original aspect ratio) so that it fills both the height and width of the display, and then is cropped (if needed) to fill the screen.

As you select different wallpaper and style options, the preview image to the right of the Background and Options selectors changes to show what your screen will look like. Once you're happy with your background, click Apply to save your changes. You can then click the Overview button at the top of this screen to return to the System Settings application's main screen.

Switching Themes

A theme is the term for the combination of window and dialog controls, a color scheme, a set of icons, and a background that gives your desktop and the windows that it displays a specific look and feel. Many themes also come with a specific desktop background to complete the visual experience that their author was trying to achieve. Experimenting with different themes can be a lot of fun, and can also be extremely useful to make it easier to read window titles and generally adjust contrast on your screen.

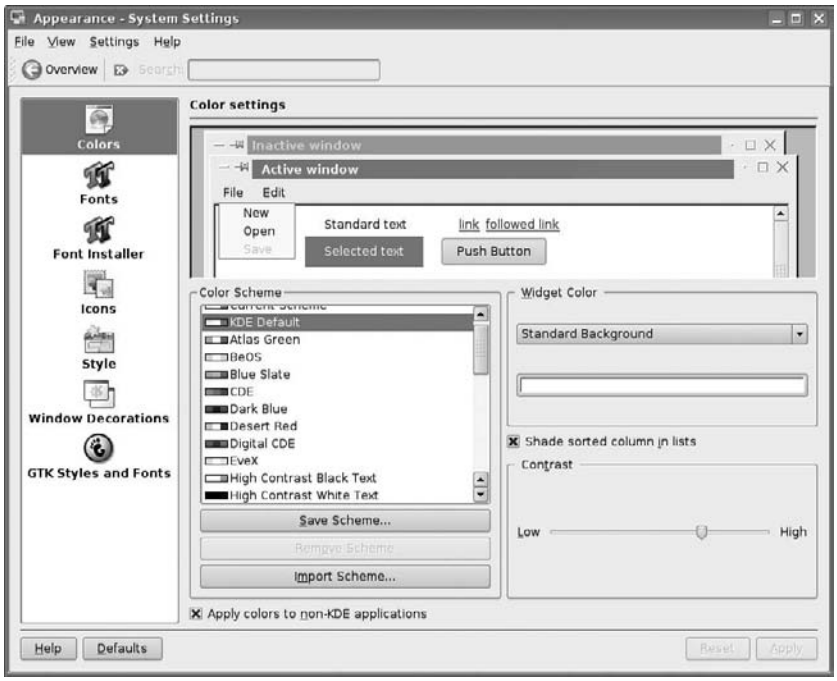
KDE 3-based Kubuntu systems provide two ways of changing the theme used on your system: you can either use the Kthememanager application to install complete themes that you have downloaded or otherwise obtained, or you can use the System Settings application's Appearance screen to configure individual aspects of your system's appearance manually. The next two sections discuss each of these and their capabilities.

Using the System Settings Application to Configure Themes

To fine-tune the theme used on your system, select the K Menu ➪ System Settings application. Once this application starts, select the Appearance icon in the Look & Feel section. The Appearance Preferences screen displays, as shown in Figure 7-34.

FIGURE 7-34

Specifying theme preferences



The Appearance screen provides the following icons that enable you to fine-tune different aspects of your desktop:

- **Colors:** Lets you configure the colors used for window borders, window and menu backgrounds, widget color, and overall contrast between foreground and background.
- **Fonts:** Sets the type and size of the fonts used by different portions of the desktop display, including toolbars, menus, window titles, the desktop itself, and so on. This screen also enables you to configure how anti-aliasing is used on your system and to force fonts to be displayed at a specific DPI.
- **Font Installer:** Enables you to manually install bitmap, OpenType, TrueType, and other external fonts for use on your system. By default, these installed fonts will only be available to the user who installed them, but they can also be installed system-wide so that everyone can see and use them.
- **Icons:** Enables you to select between and install different icon themes.
- **Style:** Enables you to configure the behavior and appearance of the widgets used on KDE windows, the effects used when maximizing or moving windows or using menus, the appearance and visual effects used in the toolbar, and so on.

- **Window Decorations:** Enables you to configure the layout and appearance of window titles, the buttons that are present in them, and so on.
- **GTK Styles and Fonts:** Sets whether GTK (i.e., GNOME-based) applications will use KDE styles, fonts, and keyboard shortcuts, standard GTK equivalents, or custom settings. This screen also enables you to apply some customizations so that Firefox and Thunderbird will look (and work) better in the KDE environment.

The System Settings ➪ Appearance screens enable you to fine-tune the theme, graphics, fonts, and almost any other aspect of your desktop. However, if you want to simply download and apply existing themes that you've found on the Net, you may be more interested in using the Kthememanager application that is described in the next section.

Using the Kthememanager Application to Install Themes

The KDE 3 desktop comes with many attractive themes. By default, KDE 3-based Kubuntu systems use a combination of a custom Kubuntu splash theme, window style and decoration theme called Plastik, an icon theme called Crystal SVG, and a standard KDE color scheme.

The previous section explained how to use the System Settings application's Appearance screens to customize the appearance of the theme that is currently being used on your Kubuntu system. However, many complete alternative themes are available on the Net that enable you to completely replace the theme that is currently being used on your system. My favorites KDE theme sites are KDE-Look.org (`www.kde-look.org/`), the KDE section of SkinBase.org (`www.skinbase.org/kde.php`), and Freshmeat's KDE themes section (`http://themes.freshmeat.net/browse/57/`). A quick Web search for "KDE themes" with your favorite search engine will turn up many more sites. Themes for different KDE versions are very different—make sure that you select the version that is designed for use with the version of KDE that is present on your system. The 3.*x* themes should work fine on a default Kubuntu system.

TIP **Many different KDE and GNOME applications have their own themes. When you're looking for complete KDE themes, most theme summaries will explain what the theme contains. If you are exploring the KDE-Look.org site, you can select the Themes/Styles entry from the content menu at left, and select the appropriate KDE version from the drop-down menu.**

You can install a completely new theme most easily by using the Kthememanager application, which is installed by default on Kubuntu systems and can be executed by selecting the K Menu ➪ Settings ➪ Lost and Found ➪ Theme Manager menu option. The Kthememanager application is actually a KDE configuration module, which you can also start from the command line by executing the `kcmshell kthememanager` command. Figure 7-35 shows the KDE Theme Manager application with the Plastik theme selected.

To switch to another of the themes provided with Kubuntu Linux, select it from the scrollable theme pane at the left of the KDE Theme Manager dialog. The theme pane at the right shows a preview of what your desktop will look like if you switch to the selected theme. (Not all themes contain preview images, but most do.) Once you've found a theme that you like, click Apply to switch your system to the new theme.

To install and use a theme from some other source, download it to your system and unpack it in a directory that you can write to. (Most themes are provided as gzipped tar archives, and can thus be unpacked using the "tar zxvf filename" command.) I like to organize downloaded themes into a single central directory to keep track of them, so I create a directory called `~/artwork`, and download and unpack them there.

Next, click Install New Theme in the KDE Theme Manager. A dialog displays that enables you to navigate to the directory that was created when you unpacked your new theme. In this directory, select the file with the extension kth (KDE THeme). Figure 7-36 shows the KDE Theme file for the Hacker Evolution theme (downloaded from the KDE-Look.org site) being selected in this dialog.

FIGURE 7-35

Using the KDE theme manager

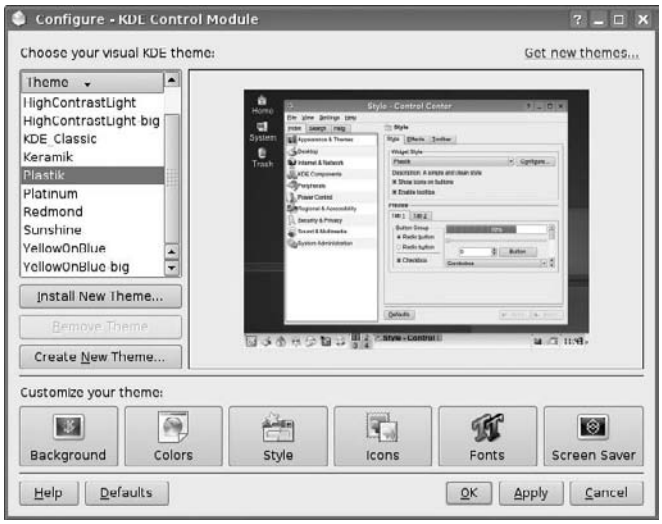
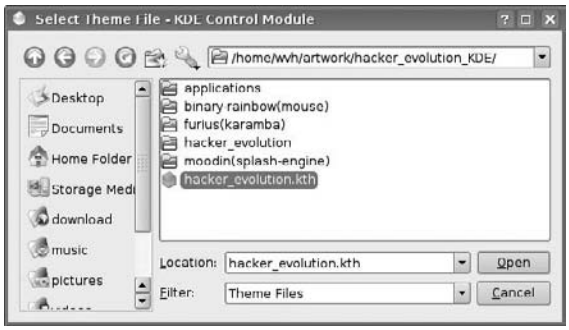


FIGURE 7-36

Selecting the theme file for a new theme



Click Open to open the theme configuration file and add the definition for the new theme to the KDE Theme Manager's left pane. If the theme contains a preview, you will be able to see a preview in the Preview pane; otherwise you will have to simply apply it and hope that you like it. if you don't like it, you can always reselect the Plastik theme to restore your basic Kubuntu theme, although you may have to adjust some of your settings using K Menu ➪ System Settings ➪ Appearance screens to restore them to their Kubuntu defaults.

Undoing KDE Configuration Changes

If you have experimented with KDE settings and have gotten your system into a state that you can't completely undo, you can always completely delete your KDE configuration settings and start over. However, this is a last-gasp solution to strange configuration settings, because it will delete all of the configuration settings that you have made for any KDE application, and should thus not be done lightly.

To remove all of your KDE settings and restore your desktop to the defaults the next time that you log in, change directory to your home directory and rename the `.kde` directory where your KDE configuration data is stored using the following command:

```
mv ~/.kde ~/.kde.BROKEN
```

You must then log out and log back in to force KDE to re-create the `.kde` directory and restore the default KDE settings. When you log back in, KDE will execute the KDE Personalizer application, a wizard that helps you set up your initial KDE configuration. Accepting all of the default values will restore your Kubuntu desktop to its default settings.

NOTE **You could always simply remove your `.kde` directory to reset your KDE environment, but renaming your bogus `.kde` directory to another directory provides you with the opportunity to examine your old settings so that you can re-create them if necessary.**

KDE Keyboard Shortcuts

"Keyboard shortcuts" is the term used to describe mapping specific commands to certain key sequences, making it easy to execute common commands by pressing a few keys at the same time rather than always having to graphically select them from menus, the panel, and so on. KDE 3-based systems provide a number of default keyboard shortcuts that are always available in various portions of the KDE desktop.

On the desktop itself, some of the most common keyboard shortcuts that are always available are as follows:

- **Alt+F2:** Displays the "run command" dialog, enabling you to enter the name of a command that you want to execute without having to start a Konsole or other terminal window. You can execute privileged commands from within this dialog by using the `kdesu` command and enclosing the entire command within quotation marks, as in the example command `kdesu kate /etc/hosts`.
- **Alt+F3:** Displays the KDE window operations menu for the currently active window.
- **Alt+F4:** Closes the currently active window without requesting confirmation.
- **Alt+F5:** Displays a list of all active windows on your system, enabling you to switch to another window by selecting it from the list.
- **Alt+Ctrl+A:** Displays any window that is currently requesting attention. This can be very useful if you have many windows open and aren't sure which one needs attention.
- **Alt+Ctrl+Delete:** Displays a confirmation dialog that asks whether you want to log out.
- **Alt+Ctrl+Shift+Delete:** Immediately logs you out of the system without asking for confirmation.

- **Alt+Ctrl+Insert:** Displays a menu that enables you to start a new session as a different user, lock the current session and start a new one, or return to any existing desktop session.
- **Alt+Ctrl+L:** Activates the screen lock application, which requires you to enter your password before redisplaying the desktop and enabling you to proceed.
- **Alt+Print:** Takes a screenshot of the current window and displays a dialog that asks where you want to save the image.
- **Alt+Tab:** Cycles through the windows on the current desktop, enabling you to quickly switch to another open application.
- **Ctrl+Print:** Takes a screenshot of the current desktop and displays a dialog that asks where you want to save the image.
- **Ctrl+Tab:** Cycles through the different desktops that are currently open, making it easy to switch desktops.

Many other preconfigured keyboard shortcuts are available, and many others can be configured. You can see all of the Keyboard shortcuts that are currently configured on your system (and define your own) by using the K Menu ➪ System Settings ➪ Keyboard & Mouse screen, and selecting the Keyboard Shortcuts icon in the left pane.

Introducing the Dolphin File Manager

A file manager is an application that enables you to navigate through the files and directories on your system and makes it easy to work with your computer's filesystem. Microsoft Windows users are probably quite familiar with Windows Explorer (and its Windows 3.1 ancestor, the aptly named File Manager), while Apple Macintosh users are similarly familiar with the Mac OS X Finder. Old-time DOS users probably even remember the Norton Commander application, which provided similar capabilities in the DOS environment, and which is still available for the Windows environment today.

If you have used KDE before, you may be surprised to learn that newer versions of KDE advocate using a file manager other than Konqueror, which has been the file browser on KDE systems since KFM (KDE File Manager) was retired long ago. Konqueror, which is discussed in the next section, has developed into a powerful tool for browsing filesystems, the Web, network locations, devices, and so on, each of which has some special needs or display characteristics that require compromises in performance, display, and so on, so that (ironically) Konqueror's power and flexibility end up limiting its growth. To solve this problem, a new dedicated file manager called Dolphin (also known as D3lphin) was introduced with later versions of KDE 3, is now shipped by default on Kubuntu 8.04 systems, and will be the default file manager for version 4 and above of KDE.

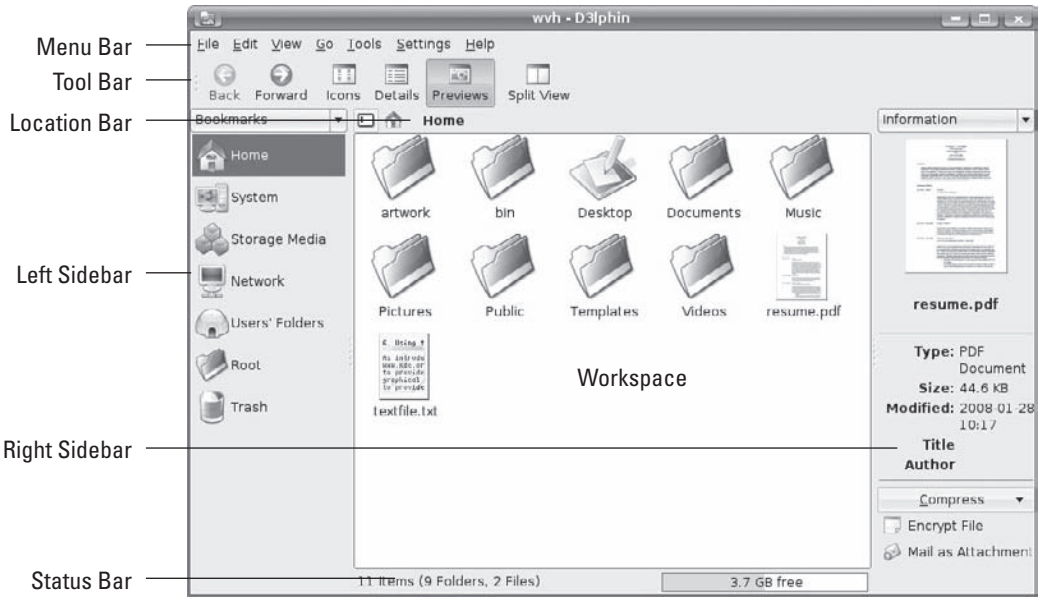
To start Dolphin, click any desktop icon that links to a filesystem location, or select K Menu ➪ System ➪ D3lphin - File Manager.

Overview of Dolphin

Figure 7-37 shows the Dolphin file manager (simply referred to as Dolphin from here on) displaying the contents of my home directory on a Kubuntu system.

FIGURE 7-37

The Dolphin file manager



Dolphin's user interface is divided into multiple sections, all of which are called out in Figure 7-37:

- **Menu bar:** Lists the menus that are available in Dolphin. These include the File, Edit, View, and Help menus that are standard in KDE applications, as well as the Go menu, which provides menu options to enable you to navigate through directory hierarchies, the Tools menu, which provided integrated access to applications such as Kfind (find files based on their names, their content, or attributes such as modified time, and so on), and the Settings menu, which enables you to configure keyboard shortcuts, Dolphin itself, and so on.
- **Toolbar:** Provides icons that simplify navigation through directories and enable you to quickly change the way in which currently selected files and directories are displayed.
- **Location bar:** Provides two icons and a menu that make it easy for you to identify your current location and quickly switch to other locations:
 - To quickly navigate to another directory by typing the full pathname of that directory, click the left-most icon over the workspace (defined below), which displays the pathname of the current directory in an editable text box. Enter the full pathname of the directory that you want to visit and press Return to display that location in the current workspace.
 - To quickly navigate to another system location in the current workspace by using the icon for your current location, left-click on the icon for your current location to see a list, and select the system location that you want to view.
 - To quickly navigate to a subdirectory of the location that you are currently viewing in a workspace, left-click on the name of the current directory and select the subdirectory that you want to visit from the drop-down menu.

- **Left sidebar:** Displays icons representing bookmarks for the default locations that are built into Dolphin, as well as any bookmarks that you have set in Dolphin. The default bookmarks that are built into Dolphin are as follows:
 - **Home:** The current user's home directory.
 - **System:** Shortcuts to common system locations, including the current user's Documents folder and Home folder, the Remote Places icon for selecting networked storage locations, the Storage Media location through which external, physical storage devices can be accessed, and the Users' Folders icon, which provides a shortcut to the directory in which users' home directories are located.
 - **Storage Media:** External storage devices that are physically connected and currently available to your system, such as hard drives, CD and DVD media, USB drives, and the like.
 - **Network:** Networked storage locations that can be accessed from your system, such as NFS or CIFS/SMB networked storage locations
 - **Users' Folders:** Centralized access to the directory where users' home directories are stored. This makes it easy to find files that you have access to in other users' directories.
 - **Root:** The root (top-level) directory of the current system.
 - **Trash:** The Trash Can for the current user, where deleted files are stored before they are purged from the system. This makes it easy to recover files that you have marked for deletion but have not yet purged by explicitly emptying the trash.
- **Workspace:** Displays icons or filenames for the directories that you are currently visiting. Whether filenames or icons are displayed in the workspace area depends on the type of view that you are currently using for the current workspace.
- **Right sidebar:** Displays information about the file or directory that is currently selected, including a list of relevant commands for the selected file or directory, and a drop-down compression menu that makes it easy for you to create a compressed version of the currently selected item.
- **Status bar:** Provides summary information about the currently selected item. If a directory is currently selected in the workspace, the status bar lists the number and type of items in that directory. If a single file is selected, the status bar lists information about file type and size. The status bar also always displays information about the amount of free space that is available on the storage device where the directory that is shown in the workspace is located.

The commands on the Settings menu enable you to configure the keyboard shortcuts used in Dolphin, the commands that are shown on the toolbar, and Dolphin itself. The aspects of Dolphin that you can configure are items such as the directory that it initially displays by default, the default view mode that it uses, and the characteristics of the various view modes, such as the size of text and icons that are displayed, and so on.

Dolphin supports three different view modes, which are selectable from the toolbar by default:

- **Details:** Displays the name, size, and last-modified date for files and directories that are currently displayed. This is much like the listings created using the `ls -l` command from a Linux command line.
- **Icons:** Displays icons for each of the files and directories that are currently displayed, using the standard icons for each type of item. A preview of the contents of known types of files is

displayed in the information sidebar (displayed at right by default) when a file is selected or when the mouse is moved over it.

- **Previews:** Displays icons for each of the files and directories that are currently displayed, but each icon for a known type of file contains a thumbnail preview of the contents of that file within the icon. As with the Icons view, a preview of the contents of known types of files is also displayed in the information sidebar (displayed at right by default) when a file is selected or when the mouse is moved over it.

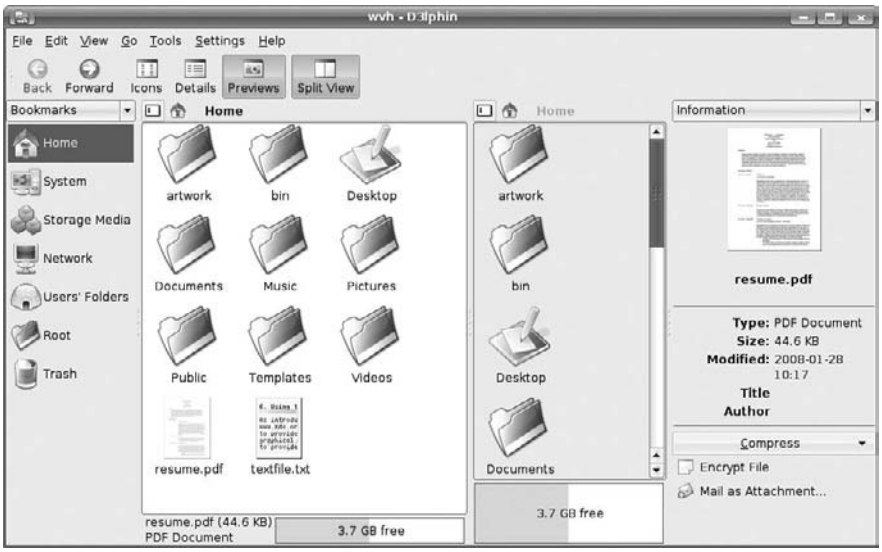
To see hidden files in Dolphin (that is, files whose names begin with a period), regardless of the view that you are using, select the View menu's Show Hidden Files command.

One of the shortcomings of many other graphical file managers is that they show only a single workspace area by default. This means that common operations such as moving files and directories either require that you have two file manager windows open (so that you can drag and drop between them), or that you have to navigate to the source location, cut files or directories, navigate to the target location, and paste those files and directories. This is tedious at best.

Both Dolphin and Konqueror provide a split, or two-pane view within a single window that makes graphically moving files and directories much easier by enabling you to navigate to the source location in one pane, navigate to the target location in the other, and then simply drag and drop files from the source workspace to the target workspace (see Figure 7-38).

FIGURE 7-38

A split view in the Dolphin file manager



Basic Operations in Dolphin

The previous sections provided an overview of Dolphin and its organization; the next few sections explain how to perform common operations to help you get started working with Dolphin if you haven't used graphical file managers before.

Selecting Multiple Files and Directories

As you work with Dolphin, you may find that you want to move or copy multiple files and directories at the same time. This is easy enough to do as long as you select all of the icons associated with those files and directories before performing a copy or cut operation. You can select multiple items in Dolphin in a variety of different ways, as follows:

- To select multiple file and directory icons, hold down the Ctrl key on your keyboard while selecting each of the icons for those files and directories.
- To select all of the files and directories that are alphabetically between two file and directory names or icons, press down the Shift key, select one filename or icon, and then select the filename or icon at the end of the range of items that you want to select. This can be tricky in KDE unless you have disabled KDE's one-click open command, as discussed in the "Customizing Mouse Behavior" section, earlier in this chapter.
- To graphically select multiple files or directories that are next to each other in a Dolphin workspace, left-click the background beside any of their icons or filenames, hold down the left mouse button, and drag the mouse cursor to display a rectangle that contains all of the icons or filenames that you want to select.

You can also combine these selection operations. For example, you could select a range of files and directories graphically by using the Shift key and then holding down the Control key to add other file or directory icons or names to the current selection.

Copying or Moving Files or Directories

To copy or move a file or directory in Dolphin, do the following:

1. Navigate to the directory that contains the file or directory that you want to copy or move.
2. Click the Split View icon to display two copies of the workspace where that file or directory is currently located.
3. In one of the two workspaces, navigate to the directory to which you want to copy or move the selected file or directory.
4. Drag the name or icon for the file or directory from the window that shows its current location to the window that shows its destination. A pop-up menu displays.
5. Select "Move Here" to move the file or directory to the new location. Select "Copy Here" to create a copy of the file or directory in the new location. You can also select "Link Here" to create a hard link for selected files, or a symbolic link for selected directories (similar to a Mac alias and a Windows shortcut).

The icon(s) for the files and directories that you had selected for copying display in the destination workspace when the copy or move operation is complete. Copying directories that contain many files may take a bit of time.

Renaming Files or Directories

To rename a file or directory in Dolphin, do the following:

1. Navigate to the directory that contains the file or directory that you want to rename.
2. Right-click on the selected icon and select Rename from the context-sensitive menu. The name of the file will be selected in Dolphin.
3. Enter the new name of the selected file or directory and press Return. The updated name displays in Dolphin.

TIP **Dolphin also enables you to rename multiple items to a new base filename. To use this capability, select multiple items before activating the rename command. The dialog that displays enables you to enter the new base name of both files, and each selected file will be renamed to the new base name with a number appended to it (e.g., "Document1," "Document2," and so on).**

Creating Directories

Creating a file is usually done by starting the application associated with some type of file and saving your work with the name of the file that you want to create. Creating directories can either be done from the command line, as explained in Chapter 14, or equally as easy from within Dolphin.

To create a directory in Dolphin, do the following:

1. Navigate to the directory where you want to create the new directory.
2. Right-click on the background of that directory, as shown in Dolphin, and select the Folder option from the Create New context-sensitive menu, or select File ➪ Create New ➪ Folder from the Dolphin menus. A dialog displays that enables you to specify the name of the new directory.
3. Enter the name of the directory that you want to create, and click OK.

Deleting Files or Directories

To delete files or directories in Dolphin, select them and right-click on the selected names or icons. Select "Move to Trash" from the context-sensitive menu. This moves the selected files and directories to the Trash can and removes them from the directory shown in the workspace, but does not actually delete them from your system.

To purge the Trash and delete all of the files and directories that it contains, select the Trash icon in the bookmarks location to display it in a workspace. You can then right-click and select the "Empty Deleted Items Folder" item from the context-sensitive menu to completely purge the Trash folder.

Getting More Information About Dolphin

Dolphin was originally a SourceForge project located at `http://sourceforge.net/projects/dolphinfileman`, but is now part of the core KDE trunk. For the latest information about Dolphin, including a discussion of its design goals, see its home page at `http://enzosworld.gmxhome.de`.

Using the Konqueror Browser

As mentioned in the previous section, Konqueror is the file manager traditionally associated with modern KDE releases. Over time, Konqueror has morphed from a file manager to an one-size-fits-all browser for many different types of objects, ranging from files to network locations to Web sites. (This is much like Microsoft's integration of Internet Explorer and the Windows Explorer, except less evil and free.) With the inclusion of the Dolphin file manager in all future KDE releases, Konqueror is still provided as a fundamental part of KDE, but is free to focus on its ability to browse Web pages and many other types of networked resources. Konqueror continues to provide a flexible framework into which plug-ins for new file types and network protocols can easily be integrated.

An icon for Konqueror is provided in the Kicker's Quick Launcher. To start Konqueror, click this icon (or select the K Menu ➪ Internet ➪ Konqueror - Web browser menu option). Figure 7-39 shows the screen that you see when you first start Konqueror.

FIGURE 7-39

The Konqueror startup screen



Konqueror is installed as the default Web browser on Kubuntu systems, and is also associated with HTML files by default. Clicking on an HTML file anywhere on your system, including inside the Dolphin file manager, will open that file in Konqueror. Figure 7-40 shows a Web site in Konqueror.

FIGURE 7-40

A Web page in Konqueror



Konqueror supports the latest version of Cascading Style Sheets (CSS 3), and supports advanced browser capabilities such as tabbed browsing, which is the ability to open multiple Web pages as tabs within a single browser window. Konqueror is generally very compliant with Internet standards. Interestingly enough, the code that does Web page rendering in Konqueror is the same code that is used by Apple's Safari Web browser on Mac OS X systems.

One of my favorite Konqueror features of all time is its support for Web shortcuts. These are special strings that enable you to go directly to a portion of a site without explicitly typing an entire address. For example, to search for something using the Google search engine, you can simply type `gg:` followed by what you want to search for in the location bar to go directly to the search results for those terms. Some other great shortcuts are:

- `dict:` *`word`*: Looks up the specified word in the online Merriam-Webster dictionary at `www.m-w.com`.
- `imdb:` *`name/movie`*: Looks up the specified actor, actress, or movie at the Internet Movie database site at `http://imdb.com`.
- `fdb:` *`word(s)/name`*: Looks up the specified words or artist in the FreeDb online music database at `www.freedb.org`

- `ths: word`: Looks up the specified word in the online Merriam-Webster thesaurus at `www.m-w.com`.
- `ucd: word(s)`: Looks up the specified word(s) in the Ubuntu community docs at `https://help.ubuntu.com/community`.
- `upkg: word(s)`: Looks for packages that contain the specified name at the Ubuntu Packages site at `http://packages.ubuntu.com`.
- `vi: word(s)`: Does a search for the specified word(s) using Vivisimo's Clusty search engine and Web site at `http://vivisimo.com`.

Many more shortcuts are available, some of which you will have to activate by checking them in Konqueror's Web Shortcuts configuration screen, which is discussed in the next section.

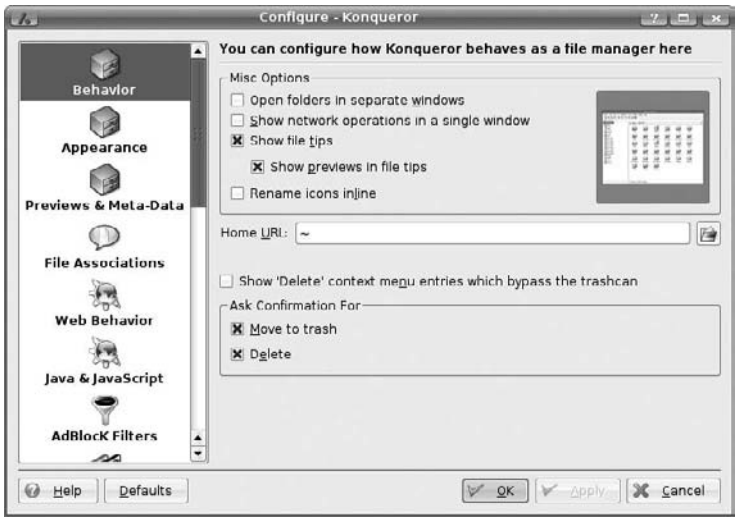
Configuring Konqueror

Konqueror provides a large number of configuration settings—too many to discuss here. This section explains how to access Konqueror's configuration settings and highlights some of the more common configuration settings that you may want to change.

To change Konqueror's default configuration, select the Settings ➪ Configure Konqueror menu option. This displays the Konqueror configuration dialog shown in Figure 7-41.

FIGURE 7-41

Configuring Konqueror



The initial Konqueror configuration screen enables you to set various characteristics of how Konqueror functions as a file manager. One of the more common of these is the "Open folders in separate windows," which creates a new window each time you select a directory, which can be very convenient if you want to move files within Konqueror by dragging and dropping rather than by cutting and pasting.

Each icon in the left portion of the Konqueror configuration screens is associated with a different configuration screen and group of related settings. Some of the more common configuration screens that you may want to use are as follows:

- **Previews & Meta-Data:** Enables you to configure the protocols for which Konqueror should be able to display file and directory previews. These include a large number of Internet protocols and local protocols for previewing attached devices and special types of files.
- **File Associations:** Enables you to configure the applications that will be used to open different types of files from within Konqueror.
- **Web Behavior:** Enables you to configure the basic characteristics of Konqueror when used as a Web browser, such as how bookmarks are saved and displayed, whether Konqueror remembers form entries and attempts to auto-complete them when new or identical forms are visited, whether to use tabs or new windows when selecting links, and so on.
- **Java & JavaScript:** Provides two tabs that enable you to configure whether Java is supported (it is, by default), and the JavaScript policies that are allowed on your system.
- **AdBlock Filters:** Enables you to define and edit patterns that Konqueror can use to recognize embedded URLs that correspond to advertisements. When a pattern is matched, Konqueror will not display or follow that URL.
- **Web Shortcuts:** Provides access to a list of predefined shortcuts to specific Web sites. See the previous section for an explanation of using Web shortcuts. This configuration screen also enables you to define your own shortcuts to Web sites (or other locations) that you frequently visit in Konqueror.

As you might expect from an application that is as powerful and flexible as Konqueror, many additional configuration options are available. After making changes to any configuration screen, you can click Apply to save your changes and continue configuring Konqueror, or click OK to save your changes and exit the Konqueror configuration dialog.

Using Konqueror as a File Manager

Because Konqueror has been used as KDE's default file manager for quite a while, it provides a number of features that are tailored towards file management. Aside from the items on the file manager configuration panel discussed in the previous section, one interesting way of using Konqueror as a file manager is to split the Konqueror window, link the two panes together, and then anchor one of them. This enables you to view a specific directory in one pane, while drilling down through the subdirectories that it contains, viewing files, and so on in the other.

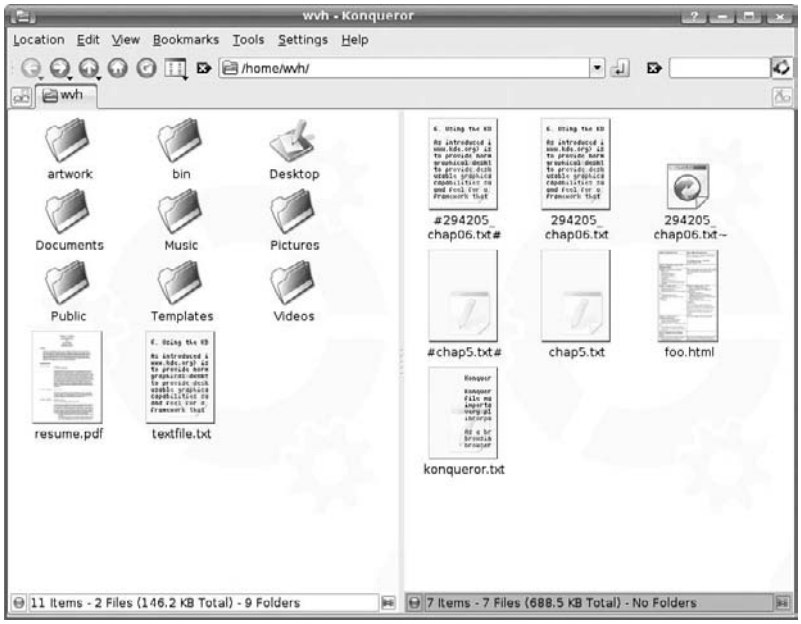
To do this, start Konqueror and use the Location ➪ Open Location menu option to specify a directory that you want to examine. Select the View ➪ Split View menu, and select either the Split View Left/Right or Split View Top/Bottom commands to create two panes, both of which show the same directory.

To link the two panes together, left-click the small box at the right of the status bar in either pane. A small icon that looks like a chain displays, indicating that the two panes are now linked together.

To configure one pane to always display the current directory, right-click on the status bar at the bottom of that pane and select Lock to Current Location. You can then click on items in the locked pane, and display the results of that selection in the right. You can also use Konqueror's back button (the left-most icon on the Konqueror icon bar) to navigate back to higher-level directories or to exit different types of file previews. Figure 7-42 shows two panes in Konqueror that have been linked in this fashion.

FIGURE 7-42

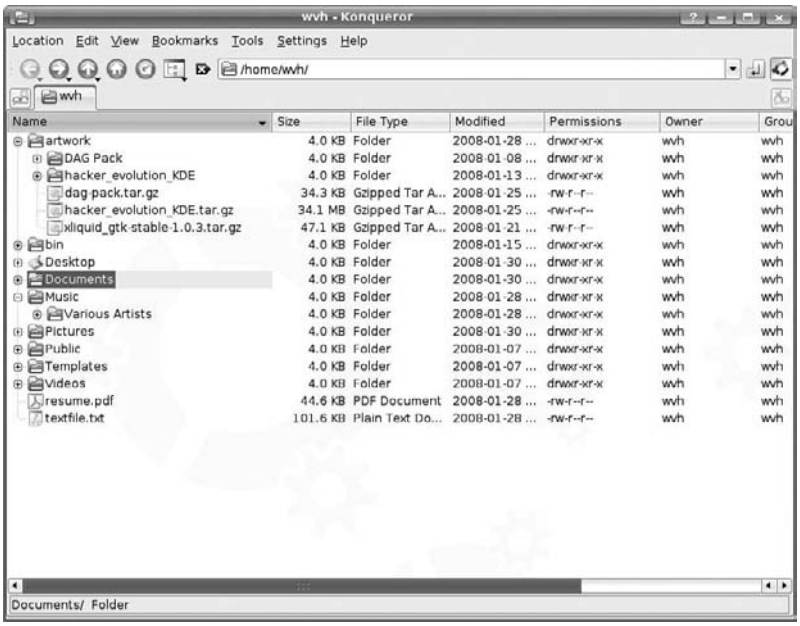
Linked panes when using Konqueror as a file manager



If you are using Konqueror as a file browser, you may want to take advantage of the different views that it supports, which are simply different ways of seeing and exploring directories. These are all listed on the View menu's View Mode submenu. I personally find the Tree View (View ➪ Tree View) mode to be quite useful when exploring various directories because it enables you to quickly and easily explore directory hierarchies. Ironically, many of the Konqueror views will start the Dolphin file manager when you click on a directory icon because Dolphin is the system-wide default for exploring directories. Konqueror's Tree View enables you to explore directories without starting Dolphin. Figure 7-43 shows Konqueror's tree view.

FIGURE 7-43

Using Konqueror's tree view



Using the Strigi Desktop Search Application

As disks and other local storage media get larger and larger, the amount of data that people save and keep around for an extended period of time gets equally large. Although this anecdote will date me, the first hard drive that I ever bought was a whopping 30MB and I remember thinking "That should last a lifetime! How can I ever fill that up?" It actually didn't take very long to fill the drive, and I remember having to clean up after myself frequently, writing files that I didn't really need to backup media so that I could free up space on my hard drive. Nowadays, some of my data files are larger than that entire drive, and the fact that my desktop system has access to several terabytes of storage mean that I'm less likely to clean up after myself anytime soon. It's simply easier to keep everything online (with offsite backups of truly critical data, of course). This is convenient when I want to consult an article that I wrote a decade ago about a defunct Linux distribution, but is not so convenient when I remember writing something a long time ago, but have no idea what file or directory it's in.

Just as search engines make it possible to find information on specific topics regardless of where it is located on the Internet, desktop search applications such as Mac OS X's Spotlight, Microsoft Windows' Desktop Search, or Google's Desktop Search make it easy to find information that is located somewhere on your desktop computer. Online search engines and desktop search applications actually work in a very similar way—they crawl the information that they are responsible for and create an index of that information against which you can submit queries.

Rather than developing yet another desktop search application, KDE 3-based Kubuntu systems use a well-known open source application called Strigi (`http://strigi.sourceforge.net/`) to index

your local storage media. Strigi can index text files, PDF documents, MP3 files, tar and zip archives, Debian and RPM packages, and OpenOffice text (`.odf`), spreadsheet (`.ods`), and presentation (`.odp`) files.

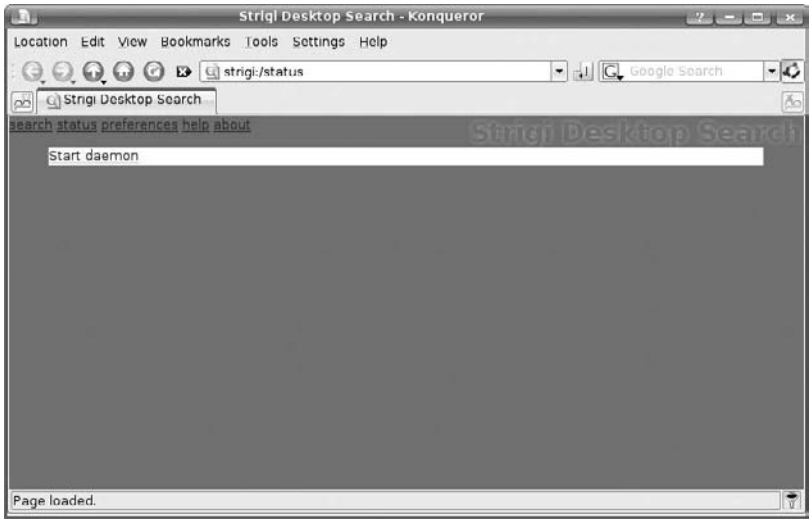
NOTE **For similar information about desktop search applications in GNOME-based Ubuntu systems, see the section of Chapter 5 entitled "Using the Tracker Desktop Search Application." That section explains how to use the tracker-search-tool which queries the index created by the tracker tool (`www.gnome.org/projects/tracker/`) on standard (i.e., GNOME-based) Ubuntu systems.**

Strigi consists of a daemon that indexes supported types of files on your local storage media, and one or more mechanisms for querying that index. The most common mechanism for querying Strigi is to select the K Menu ➪ Strigi Desktop Search menu item. A panel applet for querying Strigi is also available, and can be installed as described in the section "Customizing Panel Contents," earlier in this chapter. As an example of the power and flexibility of Konqueror, Strigi uses Konqueror for its administrative interface and to display its search results.

The first time that you start Strigi, Konqueror displays a screen like the one shown in Figure 7-44, which lets you know that the Strigi daemon is not currently running. Click Start Daemon to start the daemon, which will index all of the files in your home directory and any other directory that you set in your Strigi preferences. If an index already exists, your directories will be checked for updated files, which will automatically be re-indexed.

FIGURE 7-44

Starting Strigi for the first time



Once Strigi has finished indexing, it displays a status screen like the one shown in Figure 7-45. This screen provides some general information about the number of documents and words that were indexed, the size of the index, and so on.

At this point, you can select the search link in Strigi's Konqueror interface to display the standard query screen in Konqueror, shown in Figure 7-46.

FIGURE 7-45

The Strigi status screen

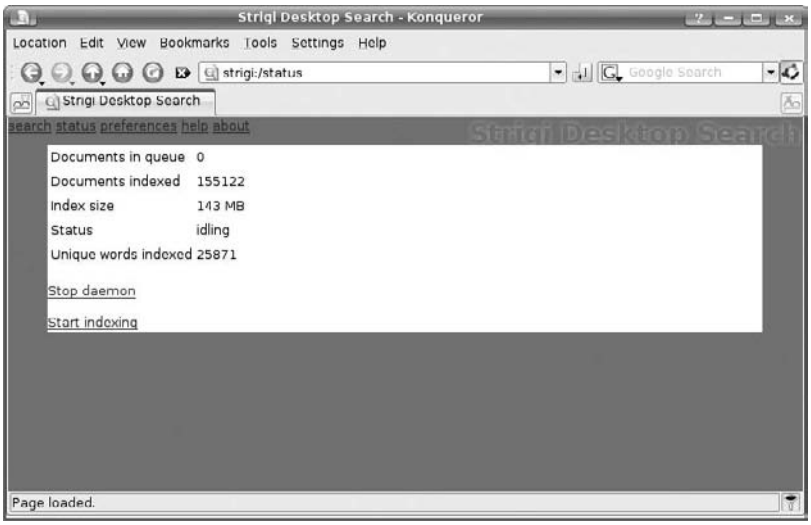
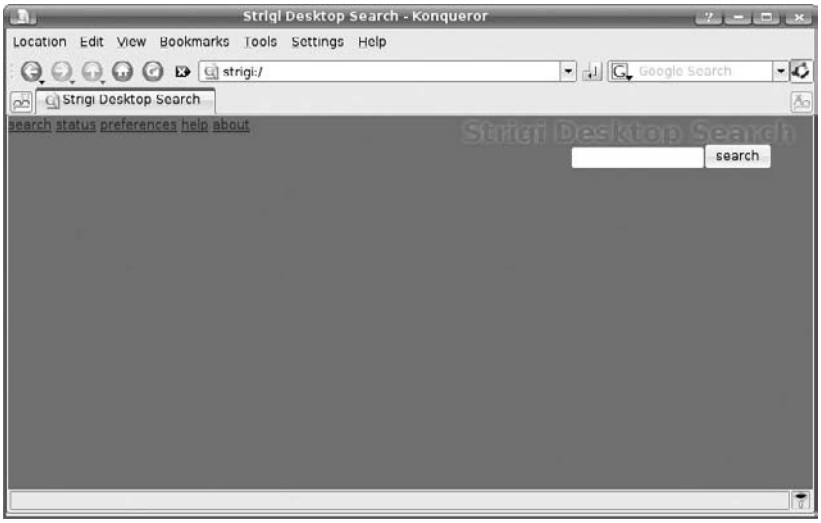


FIGURE 7-46

Strigi's query screen

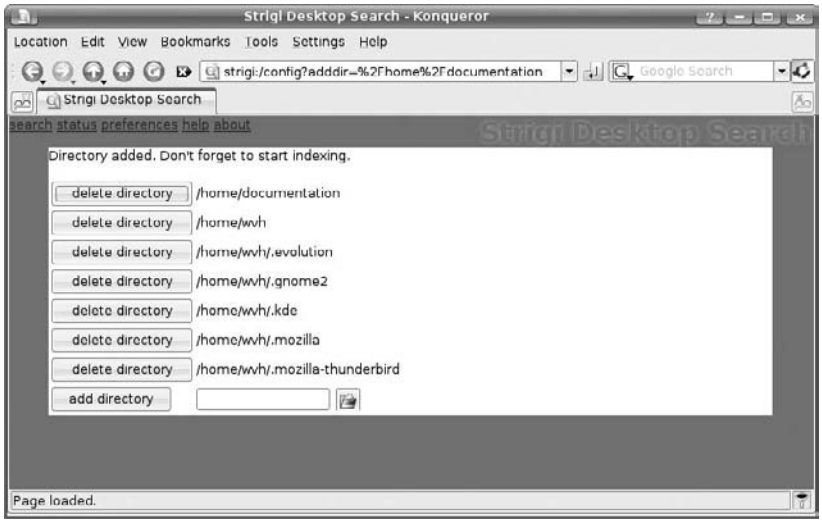


To search for a word or phrase in any of the locations that have been indexed by Strigi, enter that term in the input area at the upper right and click Search. A summary of all matching files will be displayed in Konqueror—to view any file, simply click its name or pathname, and Konqueror will open the file in whatever application is associated with files of that type.

By default, Strigi only indexes each user's home directory. To add other directories to the list of directories whose contents will be indexed, or to tell Strigi to skip certain directories, start Strigi and click the preferences link to display a screen like the one shown in Figure 7-47.

FIGURE 7-47

Strigi's preferences screen



To add another directory to the list of directories that Strigi will index, enter its name in the input area beside the Add Directory button, and then click that button. The new directory will be added to the list and a reminder message telling you to restart the indexing process will be displayed (shown in Figure 7-47). To start the indexing process, click the status link and click the Start daemon link (shown if the Strigi daemon isn't currently running) or the Start indexing link (shown if the Strigi daemon is already running).

For more information about Strigi, see its home page at SourceForge (`http://strigi.sourceforge.net`) or its Wikipedia entry (`http://en.wikipedia.org/wiki/Strigi_%28KDE%29`). Although Strigi is most commonly used in KDE environments, a GNOME panel plug-in that enables its use on GNOME systems is available as a standard DEB package (`http://packages.debian.org/unstable/gnome/deskbar-plugins-strigi`).

Changing the Window Manager for KDE

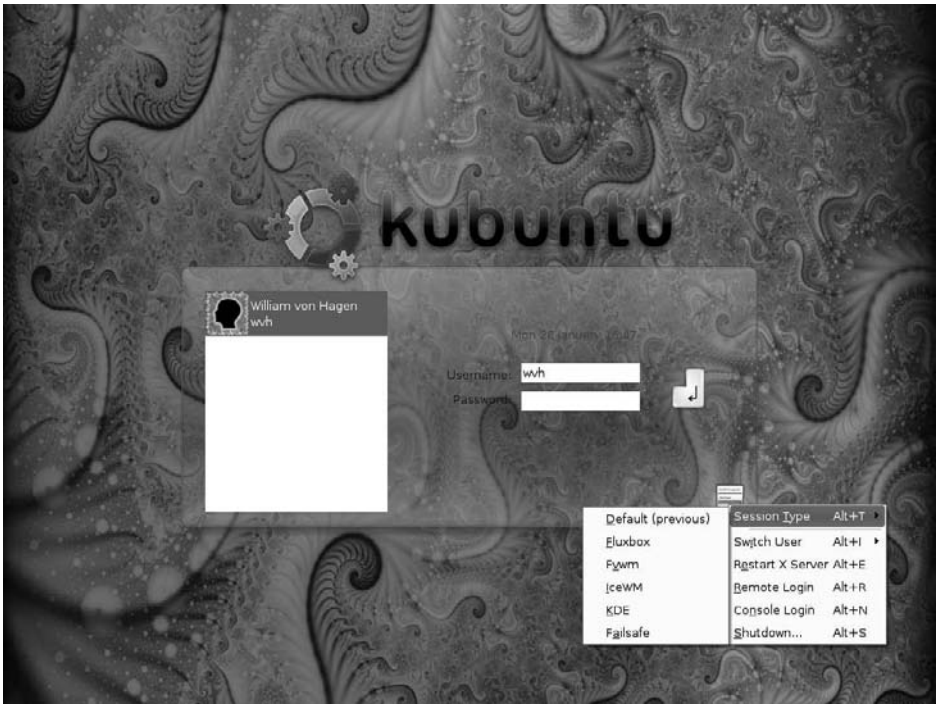
The first section of Chapter 5, "What's a Desktop? Graphical Environments for Linux," discussed the differences between window managers and desktops, and listed many of the X Window system window managers that are available today. (There are many, many more.)

After you use KDE for a while, you may be curious about what these window managers look like and how they work. Half the fun of a flexible computing environment such as Linux and the X11 Window system is experimenting with new software "because it's there." You might be especially interested in window managers if you are running Ubuntu on a slow system or a system with a limited amount of memory, such as many older laptops, because window managers require far less memory and other system resources than complete desktop environments such a GNOME and KDE do.

As you might hope, it is extremely easy to temporarily switch between using KDE and a specific window manager. First, you must enable the Universe repository, as explained in Chapter 23, "Adding, Removing, and Updating Software," because that is the place where many of the window manager packages are located. Next, you have to install one or more window managers using any of the software update tools that are also discussed in Chapter 23. After installing window managers and logging out, you can start any installed window manager by clicking the Menu button at the bottom right of the KDE display manager's login window and selecting Session Type from the menu that displays. Figure 7-48 shows a sample session menu on a Kubuntu system where a few window managers (Fvwm, IceWM, and Fluxbox) have been installed.

FIGURE 7-48

The KDM login screen's Session Type menu



Select the name of the window manager that you want to experiment with. When you log in, the selected window manager will be started for you instead of KDE.

NOTE **When experimenting with window managers, make sure that you read at least some of their documentation to ensure that their menus will be available and at least some applications can be started. It is relatively easy to start a window manager without any applications, which means that you can't do anything. This can require you to kill the window manager or the X server in order to restore control of your system.**

After experimenting with a window manager, the next time that you log out and log back in, you will need to select KDE from the Session Type menu in order to start KDE normally. If you decide that you always want to use a certain window manager, the KDE display manager will remember and automatically reuse the last window manager that you used. You can also change it permanently by creating a file called `.xession` in your home directory, which contains a command to start that window manager. For example, the IceWM window manager is installed as `/usr/bin/icewm`, so to always use that window manager, create a file called `~/.xsession` that contains the following entry:

```
exec /usr/bin/icewm
```

Make the file executable (`chmod 755 ~/.xsession`), and the next time that you log in on your system, the system will automatically start the IceWM window manager (in this example) instead of the KDE desktop. For compatibility with other ways of starting the X Window system, you should also create a symbolic link called `.xinitrc` to this file, also in your home directory. You would do this with the following command:

```
ln -s ~/.xsession ~/.xinitrc
```

Experimenting with different window managers is fun because each represents someone's idea of how you "should" interact with a graphical system, start applications, and so on. I've been using the X Window system since version 10, and thus have burned more time than I could ever calculate in experimenting with, configuring, and customizing many of the window managers that have been written in the last 20+ years. But it's been fun.

TIP **On low-memory, low-resource systems such as older laptops, you may want to start Linux without a graphical interface—if you do this, you can always start the X Window system using a command such as startx or xinit, which will also pick up the X11 startup file and symbolic link that you just created, and will therefore start your window manager for you. See Chapter 22, "Understanding the Ubuntu Startup and Shutdown Processes," for information about changing whether your system starts in graphical or command-line mode.**

The previous instructions changed whether your Kubuntu system runs KDE or a window manager. However, as mentioned earlier, all Linux desktop systems run a window manager under the covers. By default, KDE on Kubuntu systems runs the KWin window manager. If you still want to run KDE but are interested in changing the window manager that it runs underneath, this isn't hard to do. Figure 7-49 shows KDE with the IceWM window manager—I moved the KDE panel to the top to demonstrate that IceWM's standard panel is also available.

KDE uses a variety of ways to change the window manager, some more complex than others. The simplest, which is something of a hack, is to modify your `~/.xsession` file so that it first starts the window manager that you want to use, and then starts KDE. A sample `~/.xsession` file that starts the IceWM window manager, and then starts KDE is as follows:

```
/usr/bin/icewm&
exec /usr/bin/startkde
```

This is a hack because it relies on the fact that only one window manager can be active at a time. By starting IceWM before starting KDE, the KWin window manager cannot start, and thus KDE continues to run using IceWM. In general, I'm a fan of things that "just work," so this "solution" is fine with me.

FIGURE 7-49

KDE running with the IceWM window manager



NOTE **I am not recommending that you change the default window manager that is used by KDE, just pointing out how you can do so if you're in love with some specific window manager. Changing the window manager will make your Kubuntu system unlike others, and may prevent you from following some of the procedures described elsewhere in this book.**

Installing KDE on an Ubuntu System

If you are a standard GNOME-based Ubuntu user but are intrigued by the exciting capabilities and visual effects of KDE that I've discussed in this chapter, you may wish that there was a way to install and run KDE on a standard Ubuntu system. And there is! The powerful capabilities of the Debian package management system and some good package organization by the KDE developers have made it quite easy to install various "versions" of KDE and related packages by installing specific metapackages (packages that depend on other packages). The following are the different KDE-related metapackages that are available:

- **kubuntu-desktop:** If you have sufficient free space (2GB is more than enough), this package essentially does a complete KDE installation as found on a standard Kubuntu system, including all of the recommended applications such as Amarok, K3B, Kaffeine, Kontact, Kopete, and so on.
- **kde:** A slightly smaller install that only installs basic KDE packages, without the Kubuntu enhancements or recommendations.
- **kde-core:** The most minimal KDE installation, this package only installs the common KDE packages that are required to run KDE.

If you only want to install a single KDE application, using Synaptic, aptitude, or apt-get to install that application will install the KDE components that it requires (generally by installing kde-core). If you do not want to install the kubuntu-desktop package for some reason but still want to see all of the Kubuntu recommendations and enhancements, you can install the kde or kde-core packages and then install the kubuntu-default-settings package.

CAUTION **Installing the kubuntu-desktop package on an Ubuntu 8.10 or later system will install the KDE 4-based version of Kubuntu because the KDE 3-based version of Kubuntu is only supported on Ubuntu 8.04 and earlier Ubuntu releases. If you want to install the version of KDE that was discussed in this chapter on an Ubuntu system, you will have to install Ubuntu 8.04 and then install the packages in the previous list from the Ubuntu 8.04 repository.**

Summary

This chapter explored the KDE 3-based desktop that is the default user interface provided on desktop Kubuntu 8.04 and earlier systems. KDE 3 is a stable, highly usable desktop environment with a rich set of stable applications. It has been dropped from Ubuntu and Kubuntu 8.10 and later releases because KDE 4 provides a more modern, sexier envronment, but is still the choice of many people (and my recommendation for getting real work done on a Kubuntu system at the time of this writing).

After discussing the general organization of the KDE 3 desktop, the majority of this chapter focused on explaining how to customize and extend the default configuration of the KDE desktop as installed on your Kubuntu system. Personalizing your computer system's desktop is one of the first steps to making it truly yours. KDE provides the same types of customization and supports the same level of personalization as provided by other operating systems, such as Microsoft Windows or Mac OS X, it just doesn't cost as much. Is free a price?

The next chapter discusses KDE 4, the new, improved, family-size version of KDE that is used on Kubuntu 8.10 and later systems, and is truly the future of KDE.

Chapter 8

Using the KDE 4 Desktop

As mentioned in Chapter 7, "Using the KDE 3 Desktop," the K Desktop Environment is undergoing big changes with KDE 4. Kubuntu 8.10 and later ship with KDE 4. The official Kubuntu 8.04.1 LTS distribution ships with KDE 3.5.9, the last official and stable version of KDE 3.

To use or experiment with KDE 4:

- Install Kubuntu 8.10 or later.
- Use the Wubi installer on the KDE 4 CD to install KDE 4 as an alternate boot option on a Microsoft Windows system. Using the Wubi installer to install Ubuntu or Kubuntu as an alternate boot option on a Windows system is explained in Chapter 3.
- (Not Recommended) Install Kubuntu 8.04 and install KDE 4 side-by-side with KDE 3.5.9. You can then use the Session Type icon to select between them when you log in. Press Alt+T to display the Session Type menu before you log in and select the version of KDE that you want to run, and then complete the login process to start a session using that version of KDE. I do not recommend this option because the version of KDE 4 provided with Kubuntu 8.04 is not up-to-date, and will never be.

This chapter provides an overview of KDE 4, its infrastructure and components, and selected KDE 4 applications. KDE 4 itself is still under development and somewhat incomplete. It is a visually attractive and easy-to-use desktop environment that will have even traditional Mac OS X devotees all atwitter with excitement.

CAUTION **Don't mistake my enthusiasm for KDE 4 for a statement that it's ready for daily use by everyone. A major version change like this one is a huge undertaking—many KDE internals are still actively under development, and many KDE applications are still being ported. I'm not the only one who feels that way—the in-progress state of KDE 4 was the primary reason that the official version of Kubuntu 8.04 was still based on KDE 3.5.9, but all newer Kubuntu distributions will ship with KDE 4, which is visually impressive and is undeniably the future of KDE. It's worth a look and, depending on what you need to do on your system, may be suitable for your daily use.**

IN THIS CHAPTER

What's new in KDE 4

KDE 4 Desktop overview

Using widgets

Compositing in KDE 4

Customizing your desktop

What's New in KDE 4

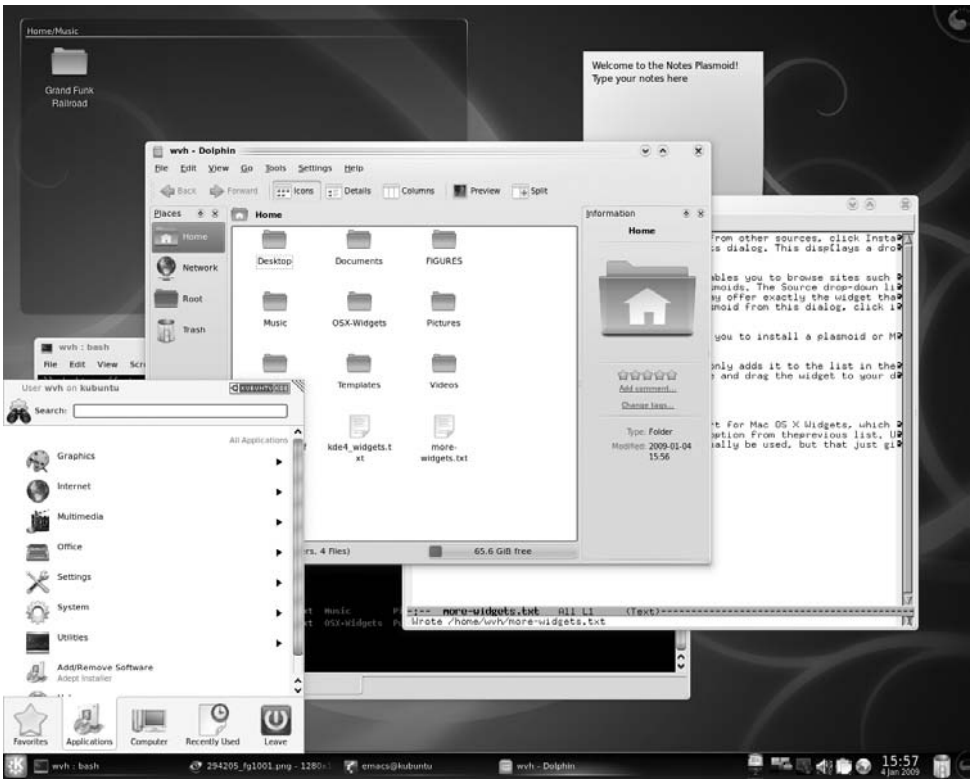
KDE 3 is a mature graphical desktop environment that has been around for a number of years, and therefore has the sorts of baggage that you associate with established projects. And there's good baggage and bad baggage—on the one hand, it is stable and mature, and any serious bugs were fixed long ago. On the other hand, it is an older codebase that implements an older model of desktop computing. The KDE 3.*x* desktop supports a fairly static set of links to things like folders, applications (via launchers), and so on. Graphically, the KDE 3.*x* desktop is primarily just a canvas on which images and X Window system applications can be displayed. Things that users typically perceive as part of the desktop, most notably the kicker panel, are actually standalone applications that have a limited amount of interaction with the desktop beyond simply being displayed over a portion of it.

Being new, KDE 4 is an attempt to address the age of the KDE 3 codebase by largely starting over and, in doing so, is able to leverage both newer auxiliary code (such as the Qt library on which it depends) and newer expectations of desktop appearance and interaction.

So what's new in KDE 4? See Figure 8-1, which shows a typical KDE 4 desktop and a few sample applications.

FIGURE 8-1

Sample KDE 4 desktop



As you can see by comparing this figure to Figure 6-1 (which shows the KDE 3 desktop), just about everything is different from a user's point of view:

- **New window decoration:** KDE 4 is based on the latest version of the Qt widget toolkit from Trolltech (now a subsidiary of Nokia), which provides the framework for developing graphical applications on KDE. Qt 4 has a completely different look and feel than previous versions of Qt, and is also much more portable than Qt 3. Rabid KDE fans are quick to point out that Qt 4 makes it much easier to port KDE 4 to other operating system environments, such as Mac OS X and Windows. While that's certainly true, I think the Cubs will win the World Series several times before a significant number of Mac users look at their computers and say, "Gee, I wish I were running KDE rather than Aqua." I love KDE and KDE 4 is beautiful, but *possibility* and *probability* are two very different words.
- **New icons and themes:** While Plasma (the new desktop shell used by KDE 4) provides a new desktop framework and Qt 4 provides a new graphical application framework, Oxygen provides new sets of icons, desktop themes, and audio effects that immediately distinguish KDE 4 from its predecessors—and in a good way. The Oxygen icons project (`www.oxygen-icons.org/`) provides a set of high-quality, photo-realistic icons in SVG (Scalable Vector Graphic) format, which inherently scale well (they are SVG, after all) and support complex 3D effects including highlighting and shadowing. The Oxygen project also provides a well-designed and well-thought-out set of standards for icon creation so that users can more quickly grasp the purpose and meaning of various types of icons. The Oxygen project follows the desktop standards established by the freedesktop.org project.
- **New desktop environment:** Plasma, which supports standard KDE applications, also provides a dynamic environment in which applets known as plasmoids execute and interact. Traditional desktop components such as the panel and panel applets are now plasmoids, which are also known as "desktop widgets," or simply widgets. Plasmoids can run on the desktop (like widgets in Mac OS X), can be implemented as containments that can contain other plasmoids (like the panel in KDE 4 and the desktop itself), or can interact with each other (by running within plasmoid containments such as the panel, like traditional applets in the Kicker). For more information about plasmoids, see the section "Using Widgets," later in this chapter. Although *plasmoids* is the Plasma-specific term for these, I'll use the more standard term *widgets* to describe these throughout this chapter. For more information about Plasma itself, see `http://plasma.kde.org`.
- **New menu/launcher organization:** The K Menu is now a widget identified as the Application Launcher, often referred to as Kickoff because that was its name when it originally appeared as a panel applet on SUSE Linux systems. This provides a very different menu experience than the traditional cascading menus found on KDE (and GNOME systems). Rather than displaying a single, cascading menu hierarchy, the new K Menu organizes available applications into a number of application groups. Selecting an application group displays only the applications in that group, and you move through the sets of available applications by selecting arrows at the left or right edge of the application launcher. For more information about the new application launcher, see the section "The K Menu, Then and Now," later in this chapter.
- **New command runner:** The traditional Alt+F2 Run Command dialog that simplified running a single command has been replaced by a search-oriented application called KRunner. It's just as easy to run a specific command as it was before, but KRunner makes it much easier to locate the "right" command. For more information about KRunner, see the section "Running Commands with KRunner," later in this section.
- **Compositing support:** KWin, the KDE window manager, has been updated to include integrated support for compositing, which makes better use of modern graphics hardware and provides sophisticated graphical effects such as shading, transparency and translucence, and so on.

(Previously, compositing for KWin was handled by a separate `xcompmgr` application.) KWin for KDE 4 will also work fine if your graphics hardware doesn't support compositing, although you won't get all of the graphical goodness provided by compositing. For detailed information about activating and using compositing in KDE 4, see the section of this chapter entitled "KWin and Compositing." For more information about compositing in general, see the beginning of Chapter 6, on the Compiz compositing window manager. For specific on KWin changes for KDE 4, see `http://techbase.kde.org/Projects/KWin`.

- **New default applications:** Beyond the porting of existing KDE 3 applications to KDE 4, many new applications are under development for KDE 4. Some of these, such as the Okular document viewer, are configured as defaults—in the case of Okular, replacing the venerable kpdf application as the default viewer for PDF and Postscript documents. Although I'm not a huge fan of file managers, the Dolphin file manager is a tasty replacement for Konqueror as the default file manager (although introduced as a default under KDE 3.*x* distributions such as Kubuntu 8.04). The Dragon media player seems to be the default for DVDs and various types of video files, and so on. This area is especially prone to change as KDE 4 matures, so the best thing I can say is that you should expect (and embrace) change. (Not a bad rule in general, actually.)

In addition to these user-visible changes, many more changes are taking place under the hood:

- **New multimedia API:** Phonon provides a single consistent API for audio and multimedia in the future, and is being integrated with existing frameworks such as gstreamer, Xine, DirectShow, QuickTime, Helix, and just about anything by writing framework-specific backends. For more information about Phonon, see `http://techbase.kde.org/Development/Architecture/KDE4/Phonon`.
- **New hardware API:** Solid provides a single consistent API for hardware, network, and power management for the future and communicates with existing frameworks such as BlueZ (the Linux Bluetooth stack), HAL (the Hardware Availability Layer), and the Network Manager through backends that can communicate with their existing APIs. For more information about Solid, see `http://techbase.kde.org/Development/Architecture/KDE4/Solid`.
- **Better support for real-time interaction:** Decibel provides a new framework for real-time communication, ranging from voice over network services such as VoIP to instant messaging. Providing a common, consistent framework makes it easier for developers to integrate advanced communication and collaboration features into their applications. For more information about Decibel, see `http://techbase.kde.org/Development/Architecture/KDE4/Decibel`.

This is only a subset of the infrastructure changes that are taking place to make KDE 4 more portable and flexible, and to better enable KDE 4 applications to interact with one another. Reusability and embedding have been a KDE trademark since KDE 2.0 thanks to its excellent KPart component technology, and you don't have to look any further than the Kontact PIM to see how well this works.

To summarize, KDE 4 looks, works, and acts better than previous KDE releases. Existing KDE users may experience a period of adjustment when moving to KDE 4, but it's a move that's well worth making once KDE 4 stabilizes and more of the cornucopia of excellent KDE applications has been completely ported to KDE 4.

The rest of this chapter provides a tour of the new KDE 4 desktop environment and many of its new features.

KDE 4 Desktop Overview

As summarized in the previous section, many aspects of the KDE 4 desktop are quite different from the way they looked and worked in KDE 3. This section discusses key aspects of using the KDE desktop, such as how to use the mouse, what various buttons do, and how to work with windows in the default theme used by KDE 4. It also covers the K menu and context-sensitive menus provided in KDE 4, and highlights the new command runner (KRunner) provided by KDE 4 to make it easier for you to find applications that may not appear on the default K menu provided in KDE 4.

Using the Mouse

Interacting with a graphical user interface of any kind traditionally requires some sort of pointing device, usually a mouse, and a keyboard to control the cursor to select and move items, access menus, and enter data within applications. As you'll see throughout this book, the actions associated with the different mouse buttons do different things depending on the portion of the screen or application in which you use them. This is known as being "context-sensitive," and is one of the keys to performing many different types of actions using only three little buttons.

Mouse actions in KDE are essentially the same as they are in KDE 3, so I'll just provide a quick review here. One big difference is that the version of KDE 4 that is available at the time of this writing does not bind any action to the middle mouse button.

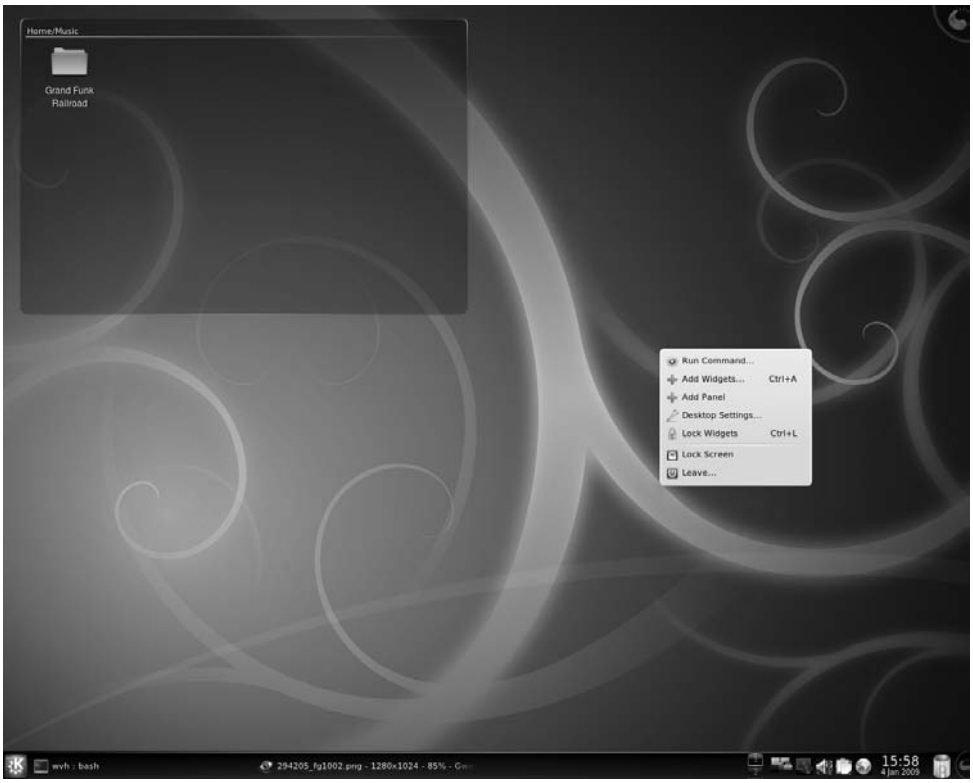
By default, the mouse in KDE 4 is configured for a right-handed user, with the different mouse buttons assigned to the following actions:

- **Left button:** Single-clicking an icon, menu, menu item, or object within an application opens the item associated with that icon. Holding down the left button and dragging the cursor in an application or on the desktop selects the icons or objects within the area that you selected. Clicking and holding down the left button on an icon and dragging that icon moves it to a new location.
- **Middle button:** No action in KDE 4. If your mouse has a scroll wheel instead of a middle button, the scroll wheel does exactly what its name suggests—enables you to scroll the contents of the current application or dialog if there is more to see.
- **Right button:** Displays any context-sensitive menu that is available at the current cursor location. On the desktop, clicking the right mouse button displays the context-sensitive menu shown in Figure 8-2. (See "Context-Sensitive Menus," later in this chapter, for more information about the desktop's context-sensitive menu.)

You can configure how responsive the mouse is, what your mouse cursor looks like, and what the various mouse buttons do by using the System Settings ➪ Keyboard and Mouse ➪ Mouse application, which is explained in Chapter 7, in the section entitled "Customizing Mouse Behavior." For example, if you are left-handed and haven't been coerced into using the mouse with your right hand, you can configure the mouse for left-handed use (which reverses the meaning of the various buttons). Other aspects of mouse behavior, such as the actions that different mouse buttons perform on the desktop, are configured in the System Settings ➪ Desktop ➪ Window Behavior application.

FIGURE 8-2

The desktop's context-sensitive menu



KDE 4 Window Decorations

It doesn't take more than a quick glance at any of the screenshots in this chapter or at your own KDE 4 desktop to notice that the default appearance of windows has changed in KDE 4. The default window decoration style (Oxygen) used by the KDE 4 desktop provides the same sorts of controls that were available in KDE 3, but things sure look different.

The controls in the title bar in the Oxygen window decoration style do the following, from right to left:

- **Title bar:** While not actually a separate control, right-clicking anywhere in the title bar for an application provides access to a menu of window-related commands that enable you to control which desktop the current window is displayed upon; move, resize, or shade (roll up) the current window; or change the behavior of the current window in relation to other windows (on the Advanced menu). Figure 8-3 shows the standard window menu in KDE 4.

 The Opacity menu, shown in Figure 8-4, is present only when special desktop effects, generally known as compositing, are enabled. These effects are discussed later in this chapter in the section entitled "KWin and Compositing." This menu lets you set how opaque the current window is—selecting a lower value makes the window more transparent, which many eye candy fans find interesting.

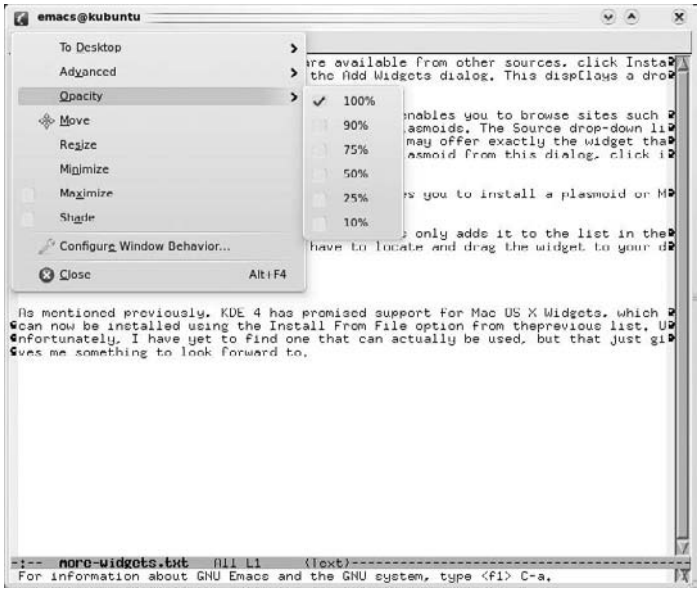
FIGURE 8-3

The window menu in KDE 4



FIGURE 8-4

Opacity settings when compositing is enabled in KDE 4



- **Minimize icon:** The down-arrow icon minimizes the current window to the taskbar. You can click that icon in the taskbar to restore the window to its previous size and screen position.
- **Maximize icon:** The up arrow makes the current window occupy the entire screen.
- **Close icon:** Closes the current window. If you've changed anything in any data file that you are working on, the application will display a dialog prompting you to save your data before closing the application. In a nice example of conceptual recursion, you can typically cancel the close operation by clicking the Close box on this dialog.

Depending on the application, you may also see a What's This? icon (a question mark inside a circle) to the left of the minimize icon. Clicking on this icon puts you in a mode where you can (left-) click on a button or other portion of the application to display a short help dialog that described the function of that button or area. Applications that provide What's This? help will also display the same information if you right-click on a button or area and select the What's This? command from the context-sensitive menu.

Menus in KDE 4

There are essentially three types of menus in KDE:

- **Context-sensitive menus** that the desktop and many applications display when you right-click on something
- **Application-specific menus** that are displayed under the title bar in any application window and which provide access to the commands, settings, and dialogs that are specific to each application
- **The K menu**, the traditional mechanism for locating and launching applications on KDE

Application-specific menus are, well, specific to each application, so there isn't much point in discussing those in a general sense. Similarly, there is no general way to discuss context-specific menus for various applications. However, both the KDE 4 desktop's context-sensitive menu (actually displayed by the KWin window manager) and the K Menu are key parts of KDE that you will certainly interact with. The next two sections discuss each of these "system" menus in more detail.

The Desktop's Context-Sensitive Menu

As mentioned in the section "Using the Mouse," context-sensitive menus are displayed when you right-click the KDE desktop, on the title bar of an application window, or in the window that is being used by an application. Figure 8-5 shows the context-sensitive menu that displays when you right-click on the KDE 4 desktop.

This menu provides the following options:

- **Run Command:** Displays the new KRunner application that enables you to locate and execute a specific command. Using this dialog is explained in the section "Running Commands with KRunner," later in this chapter.
- **Add Widgets:** Displays the Add Widgets dialog that enables you to add widgets to the KDE 4 desktop or panel. Using this dialog is explained in the section "Adding Widgets," later in this chapter.

FIGURE 8-5

The KDE 4 desktop's context-sensitive menu



- **Add Panel:** Adds a new panel to the top of your screen. See the section "Using the KDE 4 Panel," later in this chapter, for more information.
- **Desktop Settings:** Displays the dialog shown in Figure 8-6, which enables you to control the background image or slideshow that is shown on the KDE 4 desktop, whether icons can be displayed on the desktop, and whether those icons can be positioned randomly or must snap to an invisible grid to keep things neat and orderly.

 Clicking the New Wallpaper button shown in Figure 8-6 displays the dialog shown in Figure 8-7, which enables you to retrieve and install wallpaper from KDE-related sites such as the KDE-Look.org site.

FIGURE 8-6

Configuring the KDE 4 desktop



FIGURE 8-7

Adding wallpaper to the KDE 4 desktop



After installing additional wallpaper images, click Close to return to the dialog shown in Figure 8-6. You can then click the Picture drop-down to select the wallpaper that you want to use. The screen in the upper right of the dialog shown in Figure 8-6 displays a preview of the new wallpaper, and you can click Apply to display it on your desktop.

Selecting Slideshow from the Picture drop-down as the type of background enables you to identify a directory containing images that will be displayed as the background, and enables you to specify the frequency with which those images will change, as shown in Figure 8-8. A default directory of screen images, `/usr/lib/kde4/share/wallpapers`, is provided as a starting point, but you can easily add your own directory of images by clicking Add Directory and navigating to the directory that contains your images.

FIGURE 8-8

Configuring a slideshow as your KDE 4 desktop background



- **Lock Widgets:** Prevents you from moving, modifying, or adding KDE 4 desktop widgets. After selecting this option, the Add Widgets and Add Panel menu items are no longer displayed on the desktop's context-sensitive menu, and this option changes to Unlock Widgets, which you must select in order to move, modify, or add widgets.
- **Lock Screen:** Locks the screen, requiring that you specify a login session to regain access, and enter your password to do so
- **Leave:** Displays the Shutdown and Logout options dialog, shown in Figure 8-9, which enables you to log out, shut down the computer, or restart the computer

The KDE 4 desktop's context-sensitive menu provides easy access to a number of commands for controlling what is displayed on your desktop, how these items are displayed, and how you can interact with them. Most of these are things that you will set once and then not worry about again, but it's nice to know that they're all only a mouse click away.

FIGURE 8-9

The dialog displayed by the desktop menu's Leave option



The K Menu, Then and Now

The K Menu, displayed by default in the lower-left corner of any KDE screen, is probably the most important and certainly the most frequently used menu in any version of KDE. Older KDE systems displayed the K Menu (and other menus) in the Kicker panel, and KDE 4 displays it in the same location in the Panel plasmoid/widget container. However, if you are an existing KDE user, you're in for quite a surprise when you click on its icon. The K Menu is now a widget, known as the Application Launcher, which is based on the Kickoff application launcher that was originally developed on SUSE Linux systems. Because we're all familiar with the K Menu name, I'll refer to this widget as the K Menu throughout the rest of this section, even though its official name in KDE 4 is the Application Launcher. Figure 8-10 shows the new K Menu as it looks when you first left-click on its icon. (I've added a few of my favorite applications to the default choices that are shown on this menu.)

TIP **Traditional K Menu fans, don't despair! A version of the K Menu application launcher that looks just like the traditional K Menu is also available as the Application Launcher Menu widget, which you can install in the panel, as explained in the "Adding Widgets" section, later in this chapter.**

FIGURE 8-10

The new K Menu, also known as the Application Launcher widget



The new K Menu widget organizes applications in a very different way than the traditional K menu and works very differently than the traditional K Menu. Rather than displaying a single, cascading menu hierarchy, the new K Menu organizes available applications into a number of groups. Selecting a group can display menus or applications in that group, and you move through the menus of available applications by selecting arrows at the left or right edge of the application launcher. For example, selecting the Applications group at the bottom of the K Menu displays the list of menus shown in Figure 8-11.

FIGURE 8-11

The K Menu's Applications group



Selecting the arrow to the right of the System menu shown in Figure 8-11 displays the applications in that menu, as shown in Figure 8-12.

The top of the new K Menu widget displays a search field that enables you to search for applications, regardless of the group or menu location in which they are located. Once you locate an application by searching, as shown in Figure 8-13, you can click its name to execute it or right-click to display a context-sensitive menu that enables you to add it to the Favorites group of the new K Menu (shown in Figure 8-10), add an icon that serves as a launcher for that application to the desktop, or add an icon that serves as a launcher for that application to the panel itself.

FIGURE 8-12

Applications on the System menu



If an application is already located in the K Menu's Favorites group, the context-sensitive menu displays the Remove from Favorites option, rather than the Add to Favorites option.

The icons along the bottom of the K Menu widget represent the different high-level groups into which applications are organized, as follows:

- **Favorites:** Provides launchers for applications or documents that you frequently use, and is therefore the default group shown when you click the K Menu icon. As mentioned previously, you can add items to or remove them from your Favorites menu by right-clicking on them and choosing either Add to Favorites (if they are not already in the favorites group) or Remove from Favorites (if they are already in the Favorites group).
- **Applications:** Provides access to application launchers that are organized into different application groups, shown as menus. To display the items on a specific menu, click the arrow at the right side of the K Menu. To move back to a previous menu, click the large arrow at the right of the K Menu.
- **Computer:** Provides access to the System Settings application and various physical and network storage locations. If you add a removable storage device such as a USB flash drive to your system, a section entitled Removable Storage provides access to the top level of that device.

- **Recently Used:** Displays lists of applications and files that you have recently used.
- **Leave:** Provides access to Session commands that enable you to log out, lock the screen, or log in as a different user, and System commands that enable you to suspend your system to RAM memory, suspend your system to disk, shut down, or restart your computer. Figure 8-14 shows the options on this menu.

FIGURE 8-13

Searching for an application in the new K Menu



FIGURE 8-14

Options for logging out, shutting down, or suspending your system



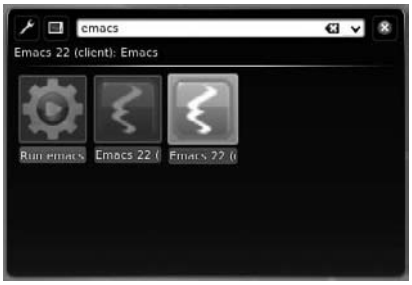
For information about the other items in the KDE 4 panel, see "Using the KDE 4 Panel," later in this chapter.

Running Commands with KRunner

The Alt+F2 keyboard command has long been a favorite way for KDE users to run a specific command without their having to fire up a Konsole window. This command is still available in KDE 4, but its functionality has been substantially improved by having been bound to an application called KRunner, which is shown in Figure 8-15.

FIGURE 8-15

KRunner, an improved way of executing a single command



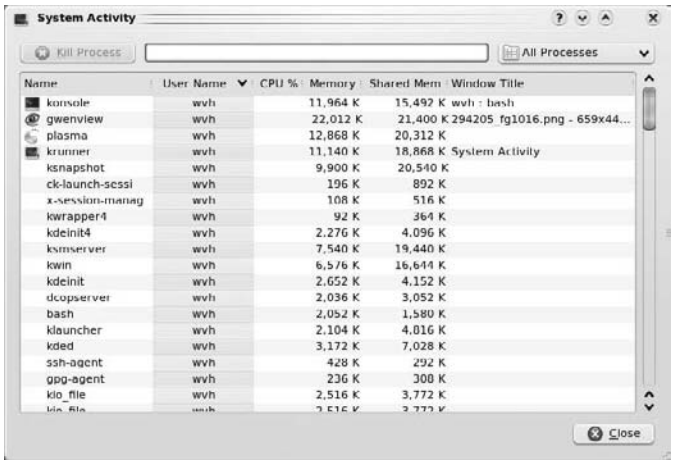
KRunner adds a search interface to the ability to run a command, and provides some nice options for how you execute a command and whom you execute it as. As you can see in Figure 8-15, typing the name of a command finds various matches for that command—in the case of my example, typing **emacs** gives you the opportunity to run the command from a virtual command line, graphically, or even running the `emacsclient` application. Clicking Show Options drops down a panel that enables you to run the command as a specific user, and even enables you to control the scheduling of the command—whether it runs as a normal process or is uninterruptible. You need root privileges to do either of these, but that's obviously not a problem on Kubuntu systems.

Another nice feature of the new KRunner dialog is the Show System Activity button, which displays a dialog that shows running processes on your system. As shown in Figure 8-16, hovering over any process in this dialog shows a tooltip that provides a meaningful description of its current state—always handy unless you inherently understand all of the columns in the output from the `ps` or `top` commands.

The Alt+F2 keyboard shortcut has always been handy for KDE users, and it's impressive to see it not only being preserved but also extended. The addition of search to this dialog is extremely convenient, and it's great to have the system activity dialog a keyboard command and mouse click away. Adding options for executing a command as a specific user is probably more useful on KDE-based distributions other than Kubuntu, but like any good bell and whistle, it's certainly better to have these options and not need them than it is to need them and not have them.

FIGURE 8-16

System activity information from KRunner



Using the KDE 4 Panel

As mentioned previously, the KDE 4 panel that replaced the traditional Kicker application is actually a special type of plasmoid, which is the Plasma-specific term for an applet that runs in the context of the new KDE 4 Plasma desktop shell. Plasmoids, more generally known as widgets, are discussed in more detail later in this chapter, but the panel is such a fundamental (and traditional) way of interacting with KDE that it's important to talk about it separately.

From left to right, the default KDE 4 panel shown at the bottom of Figure 8-1 contains the following widgets:

- **Application Launcher:** Identified by an updated version of the traditional K Menu icon, the KDE 4 application launcher was described in detail in the section entitled "The K Menu, Then and Now," earlier in this chapter.
- **Task manager:** Displays a control region for each application that is active on your current workspace. Moving the cursor over the icon for an application in the Task Manager shows a preview of the window that is currently being used by that application. Left-clicking on this control minimizes and maximizes the application, while right-clicking on this control displays a context-sensitive menu for moving and controlling that application's window(s).
- **Pager:** Enables you to manage and switch between multiple workspaces and provides a miniature display that mirrors each active workspace. One especially cool feature of the pager is the ability to drag windows from one miniature workspace view to another, which also moves the associated applications from one workspace to another.

- **System Tray:** Provides a display area, much like that found on Microsoft Windows systems, in which applications and widgets can display icons that represent system events and active services, such as receiving e-mail or alerts about low-power conditions, and so on, or provide access to system resources such as the Klipboard used for cutting and pasting data between applications, the KWallet password safe, and so on. One important application that displays messages here is the Adept Updater, which displays an icon in the notification area to let you know that system updates are available.
- **Device Notifier:** Displays a pop-up window that identifies any hardware changes to your system during the current session, such as plugging in a removable drive.
- **Digital Clock:** Displays the current time in digital format. Left-clicking this widget displays a monthly calendar. Right-clicking this widget displays a context-sensitive menu from which you can select the Configure Clock menu item to display a dialog that enables you to customize how the clock and time are displayed, your time zone, and so on.

You can remove any widget from the panel by right-clicking it and selecting the "Remove this…" menu item. See the next section for information on other available widgets and how to add them to the panel.

Using Widgets

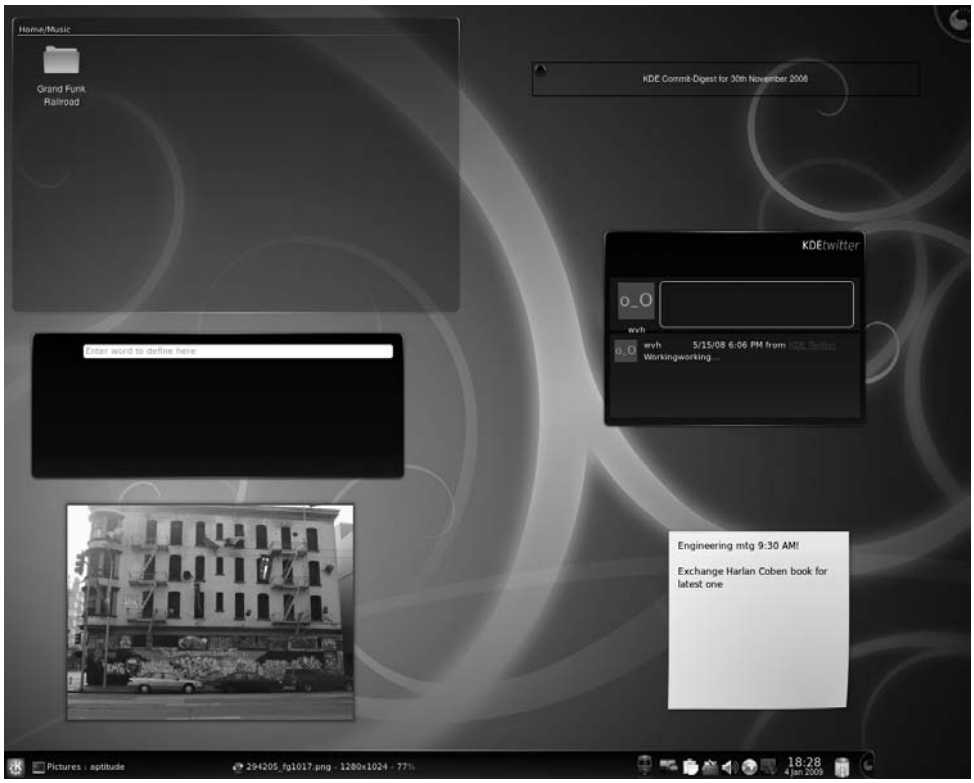
The new desktop shell used by KDE 4, known as Plasma, provides the usual necessary features, such as displaying the windows and dialogs for standard KDE applications, and also introduces a much more dynamic desktop environment in which small, graphical applications can execute. Referred to by KDE folks as plasmoids, applications of this type are generally known as desktop widgets, or simply widgets, so I'll use that term to describe them. In addition to standalone widgets that run on the desktop, plasmoids can be implemented as containments, which are simply widgets that can contain another widget. Both the panel and the KDE 4 desktop are containments. Figure 8-17 shows a KDE 4 desktop with several widgets and a panel to which other widgets have been added.

NOTE **The term "widget" can be a bit confusing because graphical interface designers have traditionally used the term "widget" to describe an element of a graphical user interface (GUI) that enables you to interact with a specific data element or object, and which serves as a low-level building block for a GUI. The term "widget," as used in this chapter, refers to a "desktop widget," which is a small graphical "helper" application that runs on the desktop of a computer system, typically providing fast access to common tasks or information such as taking notes, displaying the time and date, providing a calculator, and so on. If you're a GUI designer, I apologize for the confusion, but "widget" as used in this chapter has definitely become part of the vernacular thanks to its popularity on platforms such as Mac OS X and, now, KDE.**

Widgets are the key to big changes in KDE 4. As mentioned previously, traditional desktop components such as the panel and all of the applets that it contains are now widgets. A large number of widgets are already available, and more are under development. In addition, upcoming versions of KDE 4 will be able to run Mac OS X widgets on the KDE desktop (without having Mac OS X installed), providing KDE 4 users with instant access to hundreds of good-looking, time-tested, and Mac user–approved mini-applications.

FIGURE 8-17

KDE 4 Desktop with various widgets and updated panel



KDE 4 widgets typically separate calculation from presentation, and therefore can provide the ability to display different data in different desktop contexts, such as when being displayed on the desktop itself or when being shown in the system tray. In addition, widgets can be scaled, rotated, and moved on the desktop without interrupting their functionality.

The next section discusses how to add and manipulate widgets, the default widgets provided with KDE 4 at the time of this writing, and other sources of useful widgets to help improve your KDE 4 desktop experience. For more technical information about plasmoids, see `http://techbase.kde.org/index.php?title=Projects/Plasma/Plasmoids`.

Adding Widgets

There are several basic ways to add widgets to your KDE 4 desktop:

- Right-clicking on the desktop or panel and selecting Add Widgets from the context-sensitive menu
- Moving the mouse into the upper-right corner of the screen, left-clicking, and selecting Add Widgets from the menu that displays.

If widgets are currently locked when using either of these options, you will have to select Unlock Widgets to display the Add Widgets menu option. Selecting Add Widgets in either of these cases displays the Add Widgets dialog, shown in Figure 8-18.

FIGURE 8-18

The Add Widgets dialog



To add a widget to the desktop, select that widget in the Add Widget dialog and either click Add Widget or use the left mouse button to select that widget in the Add Widget dialog. Then hold down the left mouse button, drag the widget to the location on the desktop where you want it to display, and release the left mouse button. To add a widget to the panel, use the left mouse button to select that widget in the Add Widget dialog, hold down the left mouse button, drag the widget to the location in the panel where you want it to display, and release the left mouse button.

When a widget has been installed on the desktop, you can move it to a new location with the conventional mouse behavior—by left-clicking and holding the mouse button on the widget, dragging it to a new location, and releasing the left mouse button. At the time of this writing, you cannot move a widget in the panel, but I expect that to change in future updates of KDE 4.

Moving the mouse over a widget displays a highlighted border around that widget. (If you do not see this border, right-click on the desktop and select Unlock Widgets from the context-sensitive menu that displays.) The border around a widget provides several controls for working with that widget, as shown in Figure 8-19.

FIGURE 8-19

Controls in the border of a widget



From top to bottom, the controls on the border of a widget are the following:

- **Resize:** Left-clicking and holding this control enables you to use the mouse to resize the widget. Once the widget is the size you want, release the left mouse button.
- **Rotate:** Left-clicking and holding this control enables you to use the mouse to rotate the widget. Once the widget is rotated to the position that you want, release the left mouse button. It's not completely clear why you'd want to do this, but the option is there, just in case…
- **Configure:** Left-clicking this control displays the configuration dialog that is available for this widget. You can also configure a widget by right-clicking and selecting its configuration dialog from the context-sensitive menu that displays. This option is not available for all widgets, only those that can be configured.
- **Close:** Left-clicking this control closes the widget and removes it from the desktop. You can also close a widget by right-clicking and selecting Close This Widget from the context-sensitive menu that displays.

You can disable displaying the frame and control by right-clicking on the desktop and selecting Lock Widgets. You will not be able to resize, rotate, move, or close any widgets until you right-click on the desktop again and select the Unlock Widgets menu command.

The next two sections describe the default set of widgets provided with KDE 4 in Kubuntu 8.10, and discuss how to locate and add more widgets to the Add Widget dialog (and therefore to your system, if you so desire).

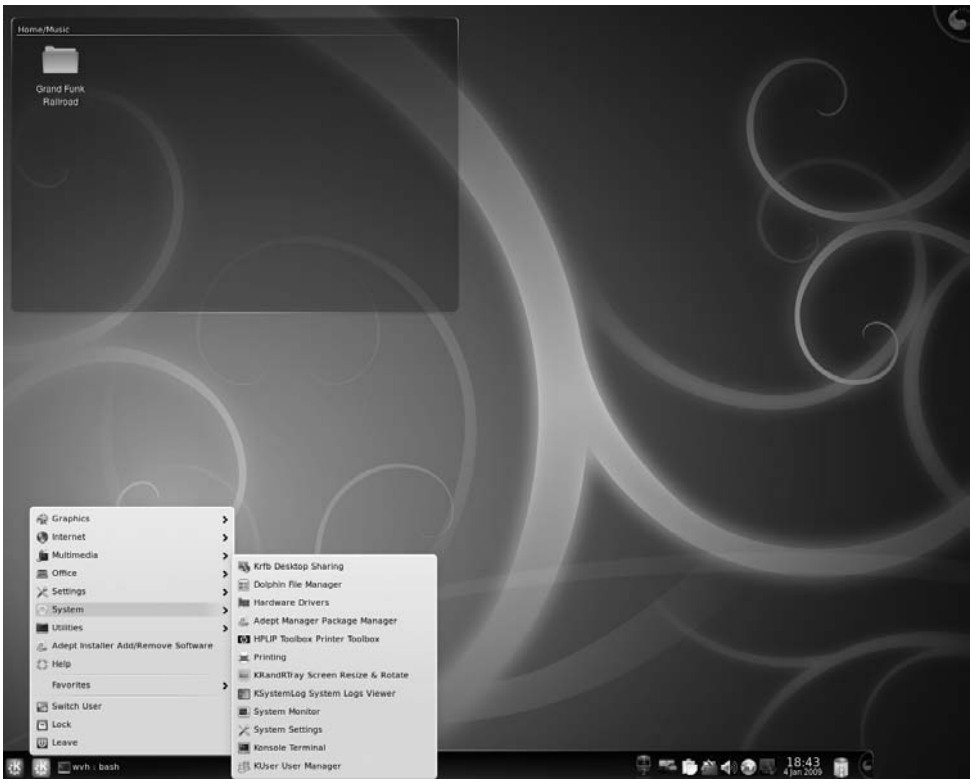
Default Widgets in KDE 4

At the time that this book was updated, a basic KDE 4 installation included the following widgets:

- **Analog Clock:** A traditional clock, with hands to indicate the time (as shown in Figure 8-19). You can click this clock to display a calendar that simplifies the lookup of past and future dates.
- **Application Launcher:** The new KDE 4 K Menu, based on SUSE's Kickoff application launcher. This widget is installed by default, at the left-most position in the panel.
- **Application Launcher Menu:** A traditional, menu-oriented KDE 3.*x* K Menu. If you are not a fan of the default Application Launcher used in KDE 4, you can drag the Application Launcher Menu widget to the panel, and then remove the default Application Launcher by right-clicking it and selecting Remove. Because this may well be one of the more popular KDE 4 widgets used by existing KDE fans, Figure 8-20 shows this widget and its menus in use in the panel of a Kubuntu system running KDE 4.

FIGURE 8-20

The traditional K Menu from the Application Launcher Menu widget



TIP **Figure 8-20 shows that you can install both the Application Launcher widget (to get the new K Menu) and the Application Launcher Menu widget (to get the traditional K Menu). This may help you prevent the withdrawal symptoms that come with a move to KDE 4.**

- **Battery Monitor:** Displays the current state and charge level for your system's battery, which is only useful on laptop or netbook computer systems.
- **Binary Clock:** Displays a Binary Coded Decimal (BCD) clock of the sort much beloved by nerds everywhere. Each column in the clock represents a digit of the current decimal time expressed in binary, with two columns each for hours, minutes, and seconds. Using this widget is much cheaper than buying an actual BCD clock, although it's not as irritating to people who are walking by your desk and want to know the current time.
- **Calculator:** Displays a simple calculator that uses direct algebraic logic, not reverse polish notation, and is therefore immediately useful for most people.
- **Color Picker:** Enables you to select and identify the color value from any portion of the screen.

- **Comic Strip:** Retrieves and displays a daily Userfriendly, Dilbert, Garfield, Peanuts, or XKCD comic strip.
- **Dictionary:** Enables you to enter a word and see its definition from the WordNet-Online dictionary (`www.wordnet-online.com`). This widget is shown in the middle left of Figure 8-17.
- **Digital Clock:** A digital clock, typically pre-installed in the lower-right corner of the panel. Like the Analog Clock widget, you can click this clock to display a calendar that simplifies the lookup of past and future dates.
- **Fifteen Puzzle:** Displays a standard 4 × 4 puzzle grid containing 15 numbers and 1 free space, where the object of the game is to arrange all 15 numbers in order—a great and classic time-waster.
- **File Watcher:** Displays changes to the end of a specified file. This is one of my favorite widgets; I use it to monitor the system messages file `/var/log/messages`. Having this info in a desktop widget makes it easy for me to spot system errors or attacks as they occur.
- **Folder View:** Displays a view of a specified folder, showing the files and other folders that it contains. An instance of this widget is displayed by default in the upper-left corner of a standard KDE 4 screen, showing the contents of your `~/Desktop` folder. The instance of this applet shown in the upper-left corner of Figure 8-17 has been customized to show the contents of my `~/Music` directory.
- **Fuzzy Clock:** Displays the current time in "fuzzy" format, meaning a semantic approximation of the current time. The configuration pane for this widget enables you to set the "Degree of Fuzziness" by dragging a slider. The settings on the slider, from left to right, correspond to five-minute increments (quarter past six, twenty past six, and so on), fifteen-minute increments (six thirty, quarter to seven, and so on), half days (Morning, Evening), or the part of the week (weekend, weekday). Like the Analog and Digital Clock widgets, you can click this widget to display a browsable calendar.
- **Kate Session Applet:** Provides a launcher for the Kate text editor that is integrated with Kate's session concept, where each Kate session is an instance of Kate with its own set of open documents, windows, toolbars, and so on. Default launcher options enable you to start Kate with no arguments (useful if you have configured Kate to always load the most recently used session), start a new Kate session, or start a new anonymous session (i.e., simply start Kate without creating a new session or loading a previous one). The Kate text editor is discussed in Chapter 15.
- **Konq Profiles Applet:** Enables you to start the Konqueror object browser using any of its standard view profiles: File Management, KDE Development (showing frequently used documentation for the KDE API, the Qt graphical development framework, the KDE bug tracking system, and the GCC GNU C Compiler, and help for the irritating "info" online help system), the text-mode Midnight Commander file manager, Tabbed browsing (showing multiple KDE-related sites), and Web Browsing (the default).
- **Konsole Profiles Applet:** Enables you to start the KDE terminal application, Konsole, using any Konsole profiles (sets of Konsole settings, tabs, and so on) that you have previously defined.
- **Lock/Logout:** A desktop button that enables you to lock the screen or display the shutdown dialog shown in Figure 8-13 with a single mouse click.
- **Luna:** Displays an applet that shows the current phase of the moon based on your time zone and configurable hemisphere.
- **New Device Notifier:** A widget that monitors your system for new hardware and displays a popup dialog when any new hardware is found. This widget is typically pre-installed in the panel.

- **Notes:** A Post-It note tool for your desktop. This widget is shown in the bottom right of Figure 8-17, just above the panel. This widget is typically pre-installed on the desktop of a KDE 4–based system.
- **Now Playing:** Provides a desktop or panel applet that displays information about any song that is currently playing in Amarok, KsCD, or other music players. This widget doesn't work correctly at the time of this writing, but let's hope that it has been fixed by the time you read this.
- **Pager:** Displays previews of the contents of all virtual desktops on your system and enables you to switch between them by clicking on the virtual desktop that you want to switch to. This widget is typically pre-installed in the panel.
- **Picture Frame:** Enables you to display a photo or other image in a frame on your desktop. You can also configure this widget to provide a slideshow of images from one or more directories, and can configure the frequency with which new images are displayed. This widget is shown at the bottom left of Figure 8-17, just above the panel.
- **Quick Access:** Displays the contents of a specified folder in list or icon mode. This widget doesn't work correctly at the time of this writing, but let's hope that it has been fixed by the time you read this.
- **Show Dashboard:** Clicking this widget switches the focus to the Plasma dashboard to simplify working with the widgets on your desktop. To return to the desktop, click anywhere on the desktop that is outside of a widget. This widget is typically pre-installed in the panel.
- **Show Desktop:** Minimizes any windows that are currently displayed on your desktop, which simplifies your work with widgets or simply enables you to hide all open windows (which can be handy if your boss walks by when you're blogging, bidding on eBay, or doing other marginally work-related tasks).
- **System Tray:** Provides a display area in which applications and widgets can display icons that represent system events and active services. This widget is typically pre-installed in the panel.
- **Task Manager:** Displays a control region for each application that is active on your current workspace, for quick application switching. This widget is typically pre-installed in the panel.
- **Trashcan:** Displays a familiar Trashcan that serves as a temporary location for deleted files. This widget is typically pre-installed in the panel.
- **Twitter Microblogging:** Enables you to receive and send tweets—obviously handy on the laptops of twitter fans. This widget is shown at the middle left of Figure 8-17.

Of course, the default set of widgets is simply a basic set of widgets that anyone might want to use. The next section explores some alternate sources of widgets that are a bit more fun, but equally useful.

Getting Other Widgets

While the basic widgets discussed in the previous section provide a fine set of system-related widgets, you'll notice that they don't include many "fun" widgets—you know, the kind that add some neat, new functionality to your desktop but are technically unnecessary. As KDE 4 continues to evolve and improve, the default widget set will certainly continue to grow, but will still focus on providing a basic set of system-related widgets.

To add additional widgets that are more fun, you can use the Adept manager, aptitude, or apt-get to add additional widgets that are available in the Ubuntu repositories, or you can use the Add Widget dialog itself to locate and install additional widgets. At the time of this writing, the only widget that is separately available in the Ubuntu archives is the News Ticker widget, which is available in the `knewsticker-`

kde4 package. This widget displays a scrolling view of news headlines from any RDF site feed (Resource Description Framework, not Reality Distortion Field). By default, this ticker shows news from the kde.org site feed, which isn't all that exciting, but you can configure it to display any new RDF feed. One of my favorites is the news feed from SlashDot, `www.slashdot.org/slashdot.rdf`. You can see this widget in action as the top-right widget in Figure 8-17.

To explore additional widgets that are available from other sources, click Install New Widgets at the lower left of the Add Widgets dialog. This displays a drop-down list with two options:

- **Download From Internet:** Displays a dialog that enables you to browse sites, such as KDE-Look.org, that provide new and updated plasmoids. The Source drop-down list lets you select between available sites that may offer exactly the widget that you're looking for. To download and add any plasmoid from this dialog, click Install to the right of its name.
- **Install From File:** Displays a dialog that enables you to install a plasmoid or Mac OS X widget from a file on your system.

Installing widgets using either of these methods only adds it to the list in the Add Widgets dialog. You will still have to locate and drag the widget to your desktop or panel in order to actually use it.

NOTE **As mentioned at the beginning of this section, KDE 4 has promised native support for Mac OS X widgets. These can now be installed using the Install From File option from the previous list. Although they can be installed, I have yet to find one that can actually be used. On the other hand, at least I still have something to look forward to.**

KWin and Compositing

As discussed in the section of Chapter 6 entitled "What is Compositing? Why All the Fuss?" compositing enables modern window managers to display sophisticated graphical effects by leveraging the power of modern video cards, off-loading various effects to video card memory, and then combining (compositing) those specially rendered images with the desktops they manage. The best-known compositing window manager today is Compiz, which is installed and enabled on standard GNOME-based Ubuntu desktop systems. Although you can use Compiz with KDE (as discussed in Chapter 7), there's no real need to do so on KDE 4 systems unless you specifically need some of the features and effects provided by Compiz. The standard KDE window manager, KWin, has been enhanced to provide compositing support, and provides its own set of compositing effects.

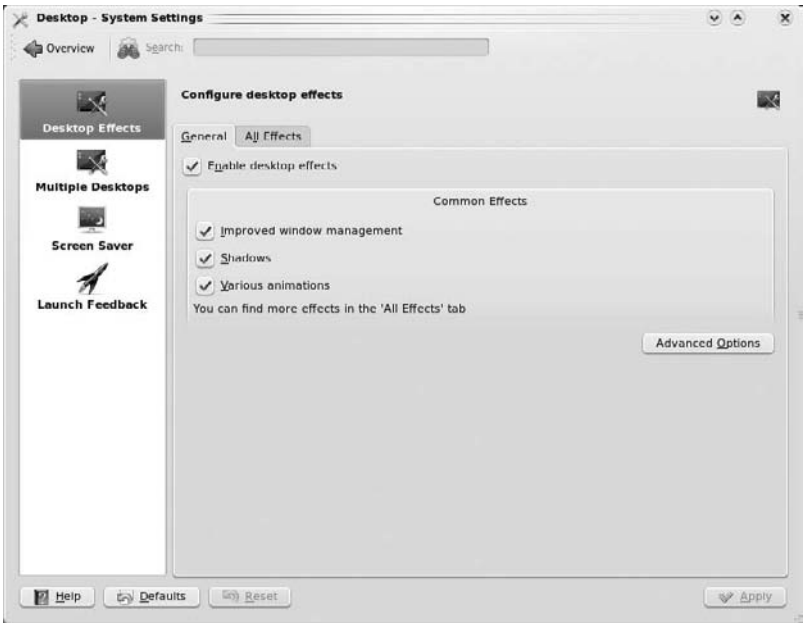
Using Compositing in KWin

Compositing support in KWin is enabled by default in the Kubuntu 8.10 release. If for some reason it is not enabled on your system, this generally means that your graphics card is not powerful enough to support the additional memory requirements of compositing. If you do not believe that this is correct or you have installed a new graphics card recently, you can enable it manually and see if it works on your graphics hardware. To do so:

1. Start the System Settings application and select the desktop icon.
2. Make sure that you're viewing the General tab of the Desktop Effects pane, as shown in Figure 8-21. (This pane should be displayed by default when you start the System Settings application.)
3. Select the Enable desktop effects checkbox, and click Apply.

FIGURE 8-21

Enabling desktop effects in KWin



Compositing effects are enabled, and a dialog with a timeout displays, asking you to confirm that you want to enable compositing support. If you do not confirm within the timeout, compositing support is disabled again.

Not all graphics cards provide sufficient memory for compositing. If your screen simply flashes on and off or the image on your desktop is distorted, you can simply wait for the desktop effects dialog to reset KWin to its standard, non-compositing settings. However, if your desktop looks the same with some immediate shadow and border effects, you're in luck, because your graphics hardware and video driver support compositing. The next section discusses the default desktop effect settings that are enabled in KWin; the following section discusses a variety of other settings and effects that you can activate to customize and fine-tune the behavior of KWin on your system.

If compositing can be enabled on your system, the Common effects portion of the Configure desktop effects dialog enables you to selectively activate or deactivate sets of compositing effects by de-selecting the Improved Window Management, Shadows, and Various Animations checkboxes. Deactivating some types of common effects can improve performance, and you may not care about them in the first place. De-selecting any of these checkboxes simply de-selects some of the desktop effects that are active by default on the All Effects tab, which is discussed in the section "Enabling and Disabling Compositing Effects," later in this chapter.

I suggest leaving the Improved Window Management option always selected, but you can try deactivating Shadows or Various Animations if you are having performance problems or simply want to try to improve graphics performance on your system. Shadows and animations are purely graphical, eye-candy effects, while the Improved Window Management effects generally provide some extra usability.

Standard Compositing Effects and Shortcuts

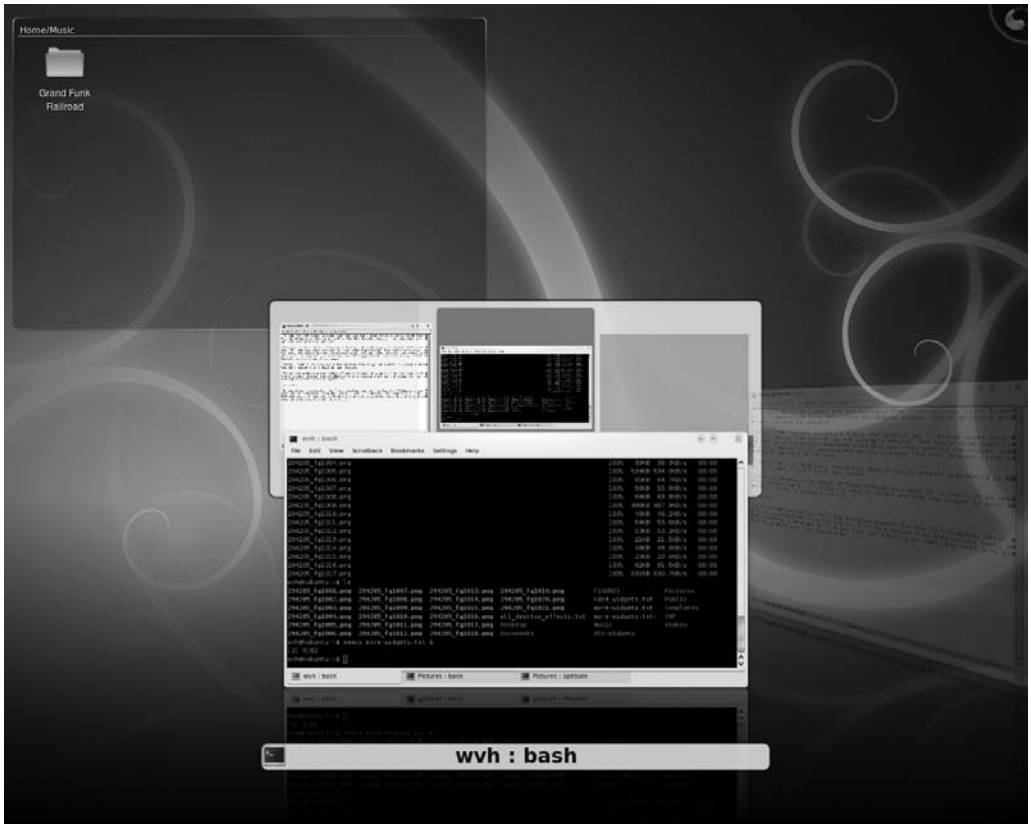
Enabling compositing in KWin activates a default set of desktop effects, many of which have associated key bindings to activate certain of their capabilities. The next section lists all of the desktop effects that are supported by KWin at the time of this writing and notes whether they are active by default. This section simply provides a quick list of keyboard commands that you can use after enabling KWin's default desktop effects.

Keyboard commands supported by KWin's default set of desktop effects are as follows:

- **Alt+Tab:** Enables you to cycle through all available windows, displaying live thumbnails of each one, as shown in Figure 8-22. To switch to a given window, release the Alt+Tab keys once the window is highlighted.
- **Ctrl+F8:** Displays all of your virtual desktops on the screen as a grid. Once this grid displays, you can select and display a specific virtual desktop by clicking on it or by using the arrow keys and pressing the Enter/Return key when the desktop that you want to switch to is highlighted.

FIGURE 8-22

Cycling through available windows in KWin

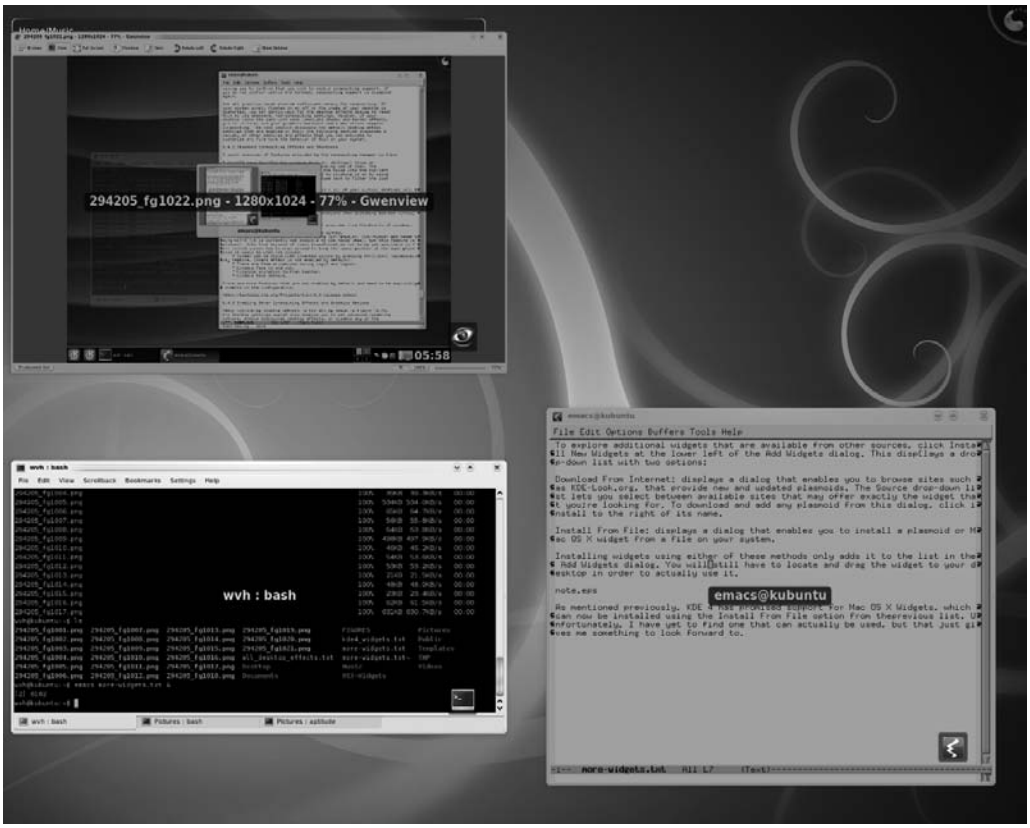


TIP **Because the grid scales all virtual desktops to fit on the screen at the same time, the grid is somewhat distorted if you have only two virtual desktops. It looks much better with four or more virtual desktops. To increase the number of virtual desktops that are available on your system, right-click on the Pager widget or the desktop, select Configure Desktop, and increase the number of virtual desktops shown at the top of this dialog. Click OK to apply your changes and close the Configure Desktops dialog.**

- **Ctrl+F9:** Displays all of the windows from a single desktop in reduced form on the screen so that you can activate a specific window. (This is similar to the Expose feature in Mac OS X.) Once the windows display, you can select one of them by clicking its image or by using the arrow keys and pressing the Enter/Return key when the window that you want to switch to is highlighted. By default, this effect can also be invoked by moving the mouse into the upper-left corner of the screen. This effect looks much like Figure 8-23, but only shows all applications from the current desktop.
- **Ctrl+F10:** Displays all of the windows from all desktops in reduced form on the screen so that you can activate a specific window, as shown in Figure 8-23. Once the windows display, you can select one of them by clicking its image or by using the arrow keys and pressing the Enter/Return key when the window that you want to switch to is highlighted.

FIGURE 8-23

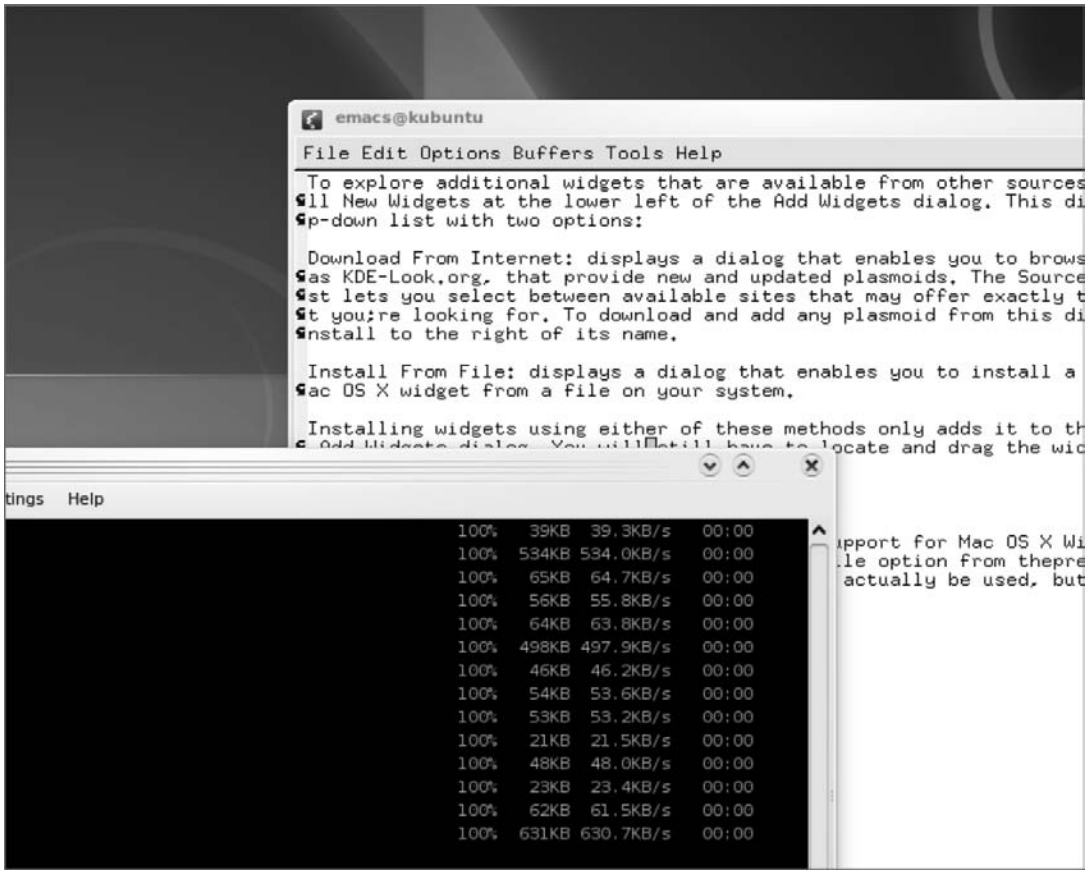
Showing all available windows in KWin



- **Meta-=/Meta—/Meta-0:** Enables you to zoom in by pressing the Meta and = keys, zoom out by pressing the Meta and - keys, or view your desktop at its standard size by pressing the Meta and 0 keys. You can change these keys by clicking the Settings button beside the Zoom effect in the Desktop Effects dialog's All Effects tab, as discussed in the section "Enabling and Disabling Compositing Effects." Figure 8-24 shows a zoomed-in view of a desktop.

FIGURE 8-24

Zooming in on a desktop in KWin



NOTE **As a joke, the Meta key is labeled with a Microsoft Windows icon on some keyboards. This key is also sometimes referred to as the Super key. The Meta/Super label is simply a traditional way of referring to a key that provides yet another way of modifying the values returned by other keys when it is pressed in combination with those keys.**

Configuring Advanced Graphics Options

After activating desktop effects in the dialog shown in Figure 8-21, the Desktop applet also enables you to set advanced rendering options, enable additional desktop effects, or disable any of the desktop effects that are enabled by default. To set advanced rendering options, click the Advanced Options button on the dialog shown in Figure 8-21. The dialog shown in Figure 8-25 displays.

FIGURE 8-25

Advanced rendering and graphics options



You can configure the following options in this dialog:

- **Compositing type:** KWin can use either OpenGL (Open Graphics Library) APIs, or the XRender (X Rendering) extension for assembling and displaying composite images. OpenGL is the default because XRender is supported in fewer X11 drivers; however, this really depends on the video card and driver that you are using.
- **Keep thumbnails of hidden windows up-to-date:** Automatically updates the thumbnails of all windows, not just those that are currently displayed. This option is disabled by default. If you have a high-performance video card with lots of memory, you may want to enable this option because it will cause thumbnails to be displayed more quickly because specific thumbnails do not have to be updated when they are displayed.
- **OpenGL mode:** Determines how and where OpenGL creates textures: Texture From Pixmap (in graphics hardware), in Shared memory, or Fallback (software-only rendering).
- **Texture filter:** Texture mapping is the process of rendering graphics on surfaces to provide 3D effects. Different texture filters define different ways of coloring *texels* ("texture pixels"), which are the basic units of surface texture maps. The different types of texture filters that are supported in KWin are as follows:

- **Nearest (fastest):** Uses the color of the texel closest to a pixel to determine the pixel color. This can result in significant anti-aliasing artifacts, such as blockiness, jaggies, and so on, but is extremely fast.
- **Bilinear:** Samples the four nearest texels and combines their colors weighted by their average distance from the pixel to determine the pixel color. This is the default texture filter used by KWin.
- **Trilinear:** Does two bilinear filter runs on different quality levels in a predefined texture map and then interpolates the results, providing exceptionally smooth color changes. This is the slowest method, but also produces the highest quality textures.

Changing the Texture filter to Nearest may improve performance considerably, but will increase the number of aliasing artifacts that you see. This is often necessary on less powerful graphics hardware.

- **Direct rendering:** Provides performance improvements by allowing direct rendering of OpenGL commands to the graphics hardware. This option is active by default, depending on your hardware.
- **Use VSync:** Waits for the vertical sync signal before displaying rendered graphics, providing much smoother graphics at the cost of performance. If this option is enabled, switching it off will increase the number of frames per second (FPS), improving performance at the expense of graphics quality.
- **Smooth scaling (slower):** Provides higher-quality graphics, but can also decrease performance, depending on the characteristics of your graphics card. This option is selectable only if the Compositing type field has been set to XRender.

You will rarely have to modify the settings on the Advanced Options dialog. If you are experiencing significant performance problems with your display, you may want to experiment with the Texture filter options, try disabling Use VSync if it is enabled, disable Keep thumbnail of hidden windows up to date if it is enabled, and so on.

Enabling and Disabling Compositing Effects

Selecting the All Effects tab displays a scrollable list of all of the graphical effects that are available in KWin, as shown in Figure 8-26.

As you can see in Figure 8-26, many of the desktop effects that are available when using compositing in KWin offer some level of configuration, which you can do by clicking the Settings button in the All Effects list. This displays a dialog such as the one shown in Figure 8-27.

One of the primary features provided by the Settings dialog for various desktop effects is the capability to change the keyboard shortcuts that activate these effects. For example, Figure 8-27 shows the Settings dialog for the Zoom desktop effect, which ordinarily uses keyboard shortcuts that use the Super/Meta key on your keyboard. If your keyboard doesn't have this key, you can use the settings dialog to define shortcuts that use keys that you actually have. To do so, click on the current keyboard shortcut for a command to display the drop-down section shown under the Actual Size entry in Figure 8-27. To change a keyboard shortcut, select the Custom radio button, click the Input button, and press the keyboard combination that you want to use. As you can see in Figure 8-27, I've changed the keyboard shortcuts for the Zoom effect to use various combinations of the Alt key and numbers because the Alt key is present on every keyboard that I own but the Super/Meta key is not.

FIGURE 8-26

All Desktop effects in KWin



FIGURE 8-27

Configuring effects using the Settings button



At the time of this writing, the list of all available desktop effects is divided into several groups of related effects:

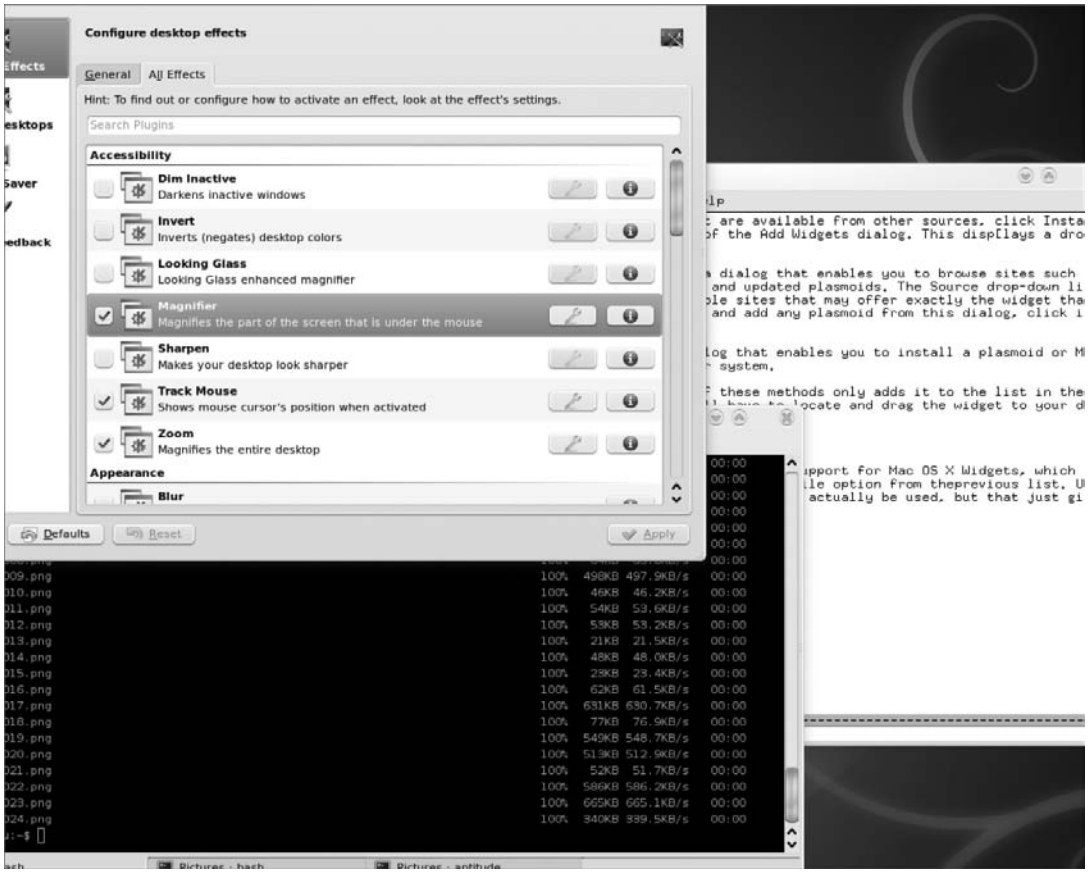
- **Accessibility:** A group of effects that make it easier to see various portions of the screen, identify active windows, and so on. The effects in this group are as follows:
 - **Dim Inactive:** Causes any inactive windows (and, optionally, the panel) to be dimmed when they are not active. This can help you locate the active window. When enabled, this option

applies only to windows. The Settings dialog enables you to extend this option to the panel and to configure just the extent to which inactive items are dimmed.

- **Invert:** Inverts the colors used on the display. I can't get this effect to work in the version of KDE 4 that is available at the time of this writing.
- **Looking Glass:** Enables magnifying selected area of the screen, much like the Magnifier effect (discussed below). I couldn't get this effect to work in the version of KDE 4 that is available at the time of this writing.
- **Magnifier:** Magnifies the portion of the screen under the current cursor location, as shown in Figure 8-28. This option shares keyboard shortcuts with the Zoom effect (shown in Figure 8-24), and will override that effect if both are enabled at the same time.

FIGURE 8-28

The Magnifier desktop effect



- **Sharpen:** Can improve the clarity of the display in certain circumstances.
- **Track Mouse:** Enables you to press the Ctrl+Meta keyboard shortcut to cause the cursor to flash, making it easy to find the current location of the mouse cursor. This option is enabled by default when desktop effects are enabled, and is part of the set of effects controlled by the Improved Window Management checkbox on the General tab.
- **Zoom:** Enables you to zoom in by pressing the Meta and = keys, zoom out by pressing the Meta and - keys, or view your desktop at its standard size by pressing the Meta and 0 keys. This effect is enabled by default.

■ **Appearance:** A group of effects that change the appearance of windows or the screen in various circumstances, providing visual clues for the state of windows or your system. The effects in this group are as follows:

- **Blur:** Blurs the background behind windows that have been configured to be translucent.
- **Dialog Parent:** Dims application windows when those applications display a dialog that requires attention. This effect is active by default, and is part of the set of effects controlled by the Improved Window Management checkbox on the General tab.
- **Dim Screen for Administrator Mode:** Dims the rest of the screen when you enter administrator mode in any System Settings applet that requires it.
- **Explosion:** Like the Fall Apart effect, this effect is supposed to provide special effects when windows close. I can't get this effect to work in the version of KDE 4 available at the time of this writing.
- **Fade:** Increases the smoothness with which windows fade in and out when they are shown or hidden. This effect is enabled by default.
- **Fall Apart:** Provides an amusing shattering effect when windows close on the screen. This is fun to experiment with, but is an effect that you may tire of quickly.
- **Login:** Provides visual effects when you log in and any default application windows display. This option is enabled by default.
- **Logout:** Dims the screen when the Logout Options dialog displays, making it easier to locate and interact with that dialog. This effect is enabled by default.
- **Minimize Animation:** Scales windows as they are minimized to the taskbar. This effect is active by default, and is part of the set of effects controlled by the Various Animations checkbox on the General tab.
- **Mouse Mark:** Draws lines on your desktop when you hold down the Shift and Meta keys and move the mouse. I can't get this effect to work in the version of KDE 4 that is available at the time of this writing.
- **Present Windows:** Enables the use of the Ctrl+F9 and Ctrl+F10 commands to display all active windows from the current desktop or from all desktops, respectively. This makes it easier to switch to another window if you have many graphical applications open simultaneously. You can use the settings dialog for this option to change the keyboard shortcuts that are associated with these effects. This effect is active by default, and is part of the set of effects controlled by the Improved Window Management checkbox on the General tab.
- **Scale In:** Provides a visual effect that expands new windows as they open on the screen.

- **Shadow:** Provides a faint shadow behind the edges of windows that is most readily visible where they overlap with others. This effect is active by default, and is part of the set of effects controlled by the Shadows checkbox on the General tab.
- **Snow:** Displays snow on your desktop, which can be entertaining for a short period of time.
- **Thumbnail Aside:** Shows a thumbnail of the current window on the edge of the current desktop. I can't get this effect to work in the version of KDE 4 available at the time of this writing.
- **Translucency:** Changes the opacity of windows, window decorations, and dialogs when they are being moved. By default, only window contents are translucent when they are being moved, although you can change this by using the Settings button. This effect is enabled by default.
- **Wobbly Windows:** Makes the edges of your windows wobble when you move a window. For some reason, this is a visual effect that I don't tire of, and can be quite useful to provide an instant check of whether compositing is enabled on a system by simply moving a window.

- **Misc:** A group of effects that don't fit into any easily defined group. The following are the effects in this group:
 - **Show FPS:** Displays a transparent window in the upper-right corner of the screen that contains a histogram that reflects KWin activities.
 - **Show Paint:** Highlights the individual portions of the screen that are being drawn and refreshed by KWin.
 - **Video Record:** Lets you use the Meta-Ctrl+V keyboard combination to toggle recording a video of the activities on your KWin desktop. This capability can be quite useful when preparing documentation or training materials.
- **Window Management:** A group of effects that control switching between and selecting windows. The effects in this group are as follows:
 - **Box Switch:** Activates an enhanced version of the Alt+Tab window switcher, which shows thumbnails of the available windows in a single box that is centrally located on the screen. This effect is enabled by default, and is part of the set of effects controlled by the Improved Window Management checkbox on the General tab. If this effect is disabled and neither the Cover Switch nor Flip Switch options is enabled, the Alt+Tab window switcher displays a simple list of available windows with no previews.
 - **Cover Switch:** Activates an enhanced version of the Alt+Tab window switcher, which cycles through a circular set of thumbnails for the available windows that are located on the current desktop. This effect is enabled by default.
 - **Desktop Grid:** Displays all of your virtual desktops on the screen as a grid, as shown in Figure 8-23. Once this grid displays, you can select and display a specific virtual desktop by clicking it or by using the arrow keys and pressing the Enter/Return key when the desktop that you want to switch to is highlighted. This effect is active by default, and is part of the set of effects controlled by the Improved Window Management checkbox on the General tab.
 - **Flip Switch:** Activates an enhanced version of the Alt+Tab window switcher, which cycles through a stack of thumbnails for the available windows on the current desktop.
 - **Taskbar Thumbnails:** Displays thumbnails of the applications associated with taskbar icons when you move the mouse over an item in the taskbar, as shown in Figure 8-29.

FIGURE 8-29

Thumbnails of applications from the taskbar



KWin's desktop effects use a plug-in architecture that makes it easy to write, distribute, and install additional desktop effects. This section discussed the desktop effects that are available in Kubuntu 8.10 at the time of this writing. Others are sure to be added in the future.

Customizing Your Desktop

Although KDE 4 delivers a completely new graphical experience with new applications, a new model for running applets on the desktop as widgets, and much more, many of the utilities used to configure different aspects of KDE 4 are essentially the same. They may look different from their KDE 3 predecessors thanks to graphical changes in KDE 4, but their capabilities, tabs, controls, and associated dialogs are almost exactly as they were in KDE 3. These types of system configuration utilities are still available through the System Settings framework, which you can access by selecting the K Menu icon in the panel, and selecting System Settings from the Favorites menu in the new KDE 4 Application Launcher, or from the top level of the KDE 4 Application Launcher Menu if you've installed that widget.

For this reason, refer to the appropriate section of Chapter 7 for information about common tasks such as customizing mouse behavior, configuring a screensaver, and so on. Other common configuration and

personalization tasks, such as changing the desktop background and activating and configuring compositing effects, have already been discussed in this chapter (in the sections "The Desktop's Context-Sensitive Menu" and "KWin and Compositing," respectively). Others have moved—for example, Power Management in KDE 4 is now a section of Display configuration, but the settings that it provides are the same.

The following sections discuss common KDE 3 desktop customization tasks that are very different in KDE 4.

Configuring Screen Resolution

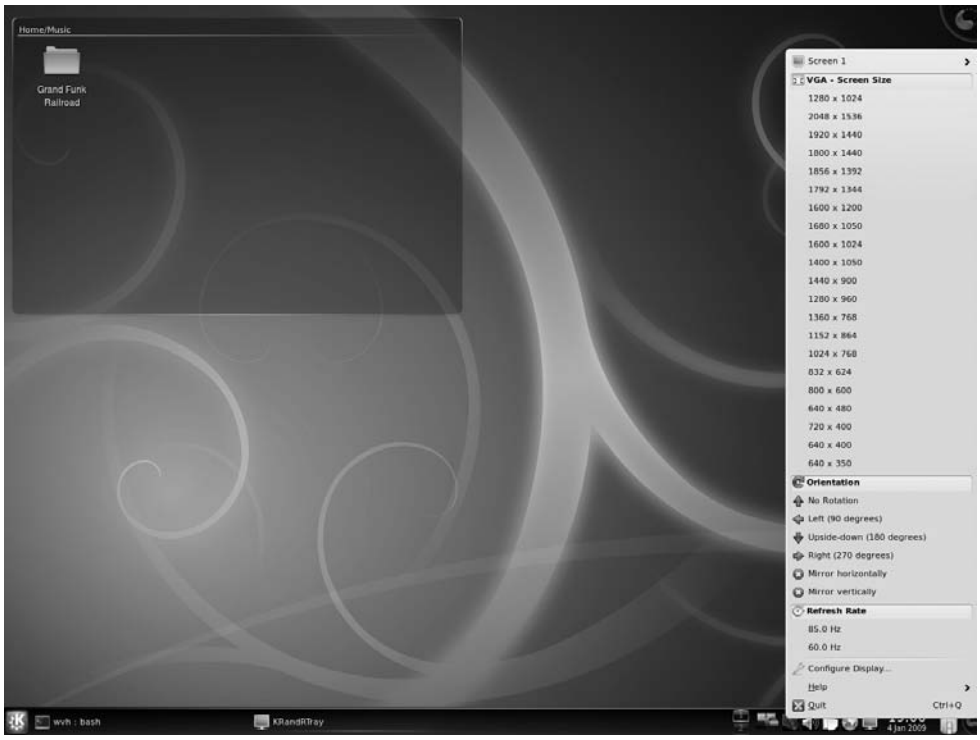
Configuring and changing display resolution settings are done in one of two different ways on KDE 4 systems:

- By selecting the K Menu's Applications section and selecting the System menu's Screen Resize & Rotate menu item
- By starting the System Settings application from the K Menu's Favorites section and selecting the Display module

The first option starts the KRandRTray application (KDE Resize and Rotate Tray) in the system tray, which you can right-click to display a menu of available screen sizes and refresh rates that you can click on to switch to a different resolution and/or refresh rate, as shown in Figure 8-30.

FIGURE 8-30

The Resize and Rotate menu from the system tray



This menu also provides a Configure Display option, which starts the standard System Settings ➪ Display module, which is essentially the second of the two methods for changing screen resolution in the previous list. Figure 8-31 shows the Display module in the context of the System Settings framework.

FIGURE 8-31

The System Settings framework's Display module



To use the Screen Resize & Rotate application from the system tray, right-click its icon in the Panel's System Tray section and select the new resolution that you want to try from the pop-up menu that displays. A dialog with a timeout displays, asking you to accept the new settings (if your screen is fine at the new resolution) or return to the previous settings. If you do not accept the new settings within the timeout, your screen automatically returns to the previous settings (which is handy if the new settings have totally screwed up your display).

The name of the Screen Resize and Rotate Tray application is somewhat misleading because you will have to use the Configure a Display menu option to start the Display setting module before you can experiment with rotation.

Using the new System Settings ➪ Display module is very similar. From the menu shown in Figure 8-30 or the dialog shown in Figure 8-31, select the new resolution and, optionally, new refresh rate that you want to try. The same sort of dialog with a timeout displays, asking you to accept the new settings (if your screen is fine at the new resolution) or return to the previous settings. As with the system tray item, if you do not accept the new settings within the timeout, your screen automatically returns to the previous settings.

Once you finish changing screen resolution and refresh rates, you can exit from the System Settings framework by selecting File ➪ Quit, or exit from the KDE Screen Resize and Rotate Tray application by right-clicking its icon and selecting Quit from the context-sensitive menu.

Switching and Creating Themes

Ordinarily (and, presumably, in the near future), you can install, apply, and create themes using the System Settings framework's Appearance module, which provides a Theme Manager module. Unfortunately, this module has not been updated for KDE 4 at the time of this writing, so themes must be installed manually at the moment.

To get new themes, go to `www.kde-look.org` and select Plasma Themes from the left-hand menu. This displays a list of available themes for KDE 4, which you can download by clicking on them and searching for the Download link.

Once you've downloaded a theme, you can install them in your own directory for your own use, or in a system directory for use by all users of your system. Themes for your own use are installed in the directory `~/.kde4/share/apps/desktoptheme`, which you will probably have to create. Themes for system use are installed into the directory `/usr/lib/kde4/share/kde4/apps/desktoptheme`.

KDE themes are packaged as gzipped tar files. To extract a theme from a downloaded theme file, change directory to one of the directories listed in the previous paragraph and extract the them using the command `tar zxvf` *`theme-file`*, where *`theme-file`* is the name of the file that you downloaded. This command will list the names of the files that it is extracting from the tar file. Make note of the top-level directory that is being created—you'll need to know that later. You will also have to preface this command with sudo if you are installing themes into the directory `/usr/lib/kde4/share/kde4/apps/desktoptheme`, because this is a system directory that is owned by the root user.

Once the theme that you want to use is installed, `cd` to the directory `~/.kde4/share/config` and edit the file `plasmarc` using a text editor. Change the line `name=default`, located after the `[Theme]` section header, to `name=theme`, where `theme` is the name of the theme that you installed. This is the name of the directory that was created when you extracted the theme from the compressed tar file.

You can then log out and log back in, at which point you will be using the new theme. You can also kill and restart Plasma from a command line by executing a command such as `killall plasma && plasma &`, but this command will produce significant output in your command-line window as you subsequently execute Plasma-related tasks such as selecting tray or panel icons, starting widgets, and so on. On the other hand, this is the easiest way to quickly see a new theme so that you can decide whether you really want to use it. You can always change the entry in the `plasmarc` file back to `name=default` to restore Plasma's default Oxygen theme.

This manual theme installation and activation mechanism is tedious, but is also short-lived. Using a System Settings module to install themes graphically will certainly be supported in updated releases of KDE 4.

Summary

KDE 4 is a beautiful work in progress that you can already use as your daily desktop if you're willing to put up with living on the bleeding edge of desktop development. This chapter highlighted the most significant changes that existing KDE users will encounter when moving to KDE 4, explained how to use the new desktop, panel, and widgets, and discussed how to activate and use new features such as the compositing capabilities that are now available in the KWin window manager. Although I use GNOME as the desktop on most of my Ubuntu systems, I have always been a KDE fan because I spent a number of years with SUSE Linux, where KDE was king. KDE 3 support on the standard Kubuntu 8.04 releases is great, and KDE 4 looks even more promising. Stay tuned!

Now that you are one with the various desktop environments provided by Ubuntu and Kubuntu, the next few chapters move into application territory by explaining how to do standard tasks such as reading e-mail (with Evolution on Ubuntu systems or Kontact on Kubuntu systems), and surfing the Web (with Firefox anywhere, or using Konqueror on Kubuntu systems).

Chapter 9

Managing E-mail and Personal Information with Evolution

After Web browsers, e-mail is the Internet's killer application. Like Web browsers, there are many different client applications for sending and receiving e-mail, especially in the open source community. In addition (and unfortunately unlike Web browsing, which basically uses HTML everywhere) there are many different protocols used to send e-mail. Your choice of an e-mail client can therefore depend to some extent on whether it supports the protocols that your Internet service provider (ISP), business, or academic institution uses to send, store, and deliver e-mail.

The easiest solution to the "which protocol, which client" problem is to use a single e-mail client that speaks every modern e-mail protocol known to humanity. Ubuntu includes an e-mail client known as *Evolution*, which lives up to this promise. Evolution is a GNOME-based mailer that was originally developed by a company called Ximian. Oddly enough, Ximian was eventually acquired by Novell, whose SUSE Linux product is largely KDE-centric, but it's at least clear that Novell knows good software when they see it.

Evolution can receive mail from mail servers that speak all of the popular e-mail protocols including POP, IMAP, IMAP4rev1, Microsoft Exchange, Novell GroupWise, and a few that would require a Google search to figure out who uses them (and why). For outgoing mail, Evolution can communicate with any Simple Mail Transport Protocol (SMTP) server or any server that is compatible with the standard Linux/Unix sendmail application.

Beyond all this, Evolution is much more than just a multiprotocol e-mail client—it is actually more of a Personal Information Management (PIM) package, of which sending and receiving e-mail is just one aspect. Evolution also provides integrated support for PIM capabilities including the following:

- Managing contacts stored locally or on Lightweight Directory Access Protocol (LDAP) or Microsoft Exchange servers
- Managing task lists stored locally or on a Microsoft Exchange server
- Managing local calendars, Web-based calendars, or Microsoft Exchange calendars

IN THIS CHAPTER

Setting up your e-mail client

Sending and receiving mail

Organizing your e-mail

Filing e-mail

Dealing with spam

Accessing webmail

Managing contacts

Managing your calendar

By now, you may see something of a thread in this introduction—compatibility with Microsoft Exchange servers and the services that they provide. Like it or not, many businesses depend on Microsoft Exchange to satisfy groupware requirements such as scheduling meetings, checking and sharing coworkers' calendars, and so on. The e-mail server portion of an Exchange Server is just a special type of IMAP mail server, but the calendaring and scheduling portions of Exchange are largely Microsoft-only. Evolution's support for complete interoperation with a Microsoft Exchange environment is truly impressive.

This chapter explains how to configure and use Evolution as a mail client, contact manager, and calendaring system. The chapter begins by discussing how to create your initial mail accounts and how to read, send, organize, and filter mail, and it discusses the powerful virtual folders feature that is unique to Evolution. It explains how to use Evolution to access popular webmail systems such as Gmail, Yahoo!, and Microsoft Hotmail/Live. It then discusses how to configure and use the contact and calendar management portions of Evolution. For even more information on Evolution's capabilities, see the section "Additional Sources of Information About Evolution" at the end of this chapter.

Starting Evolution

You can start Evolution in a variety of different ways:

- By clicking the Evolution icon in the GNOME toolbar at the top of the screen, shown in Figure 9-1. This is the easiest and most common way to start Evolution.
- By selecting Evolution browser from the Applications ➪ Internet menu.
- By typing the evolution command in a GNOME Terminal or xterm window and pressing Return.

FIGURE 9-1

The Evolution icon on your GNOME desktop



Ubuntu's GNOME desktop is also preconfigured to start Evolution for you if you click on a Uniform Resource Locater (URL) that is designed to send an e-mail message (known as a *mailto* URL) in a Web browser or other Internet-aware application.

TIP **This chapter focuses on using Evolution with the standard GNOME-based version of Ubuntu. You can also use Evolution if you are running Kubuntu by installing the evolution package (which will also install all of the GNOME requirements for running Evolution), although Kubuntu already includes an excellent PIM package called Kontact. Kontact is discussed in Chapter 10.**

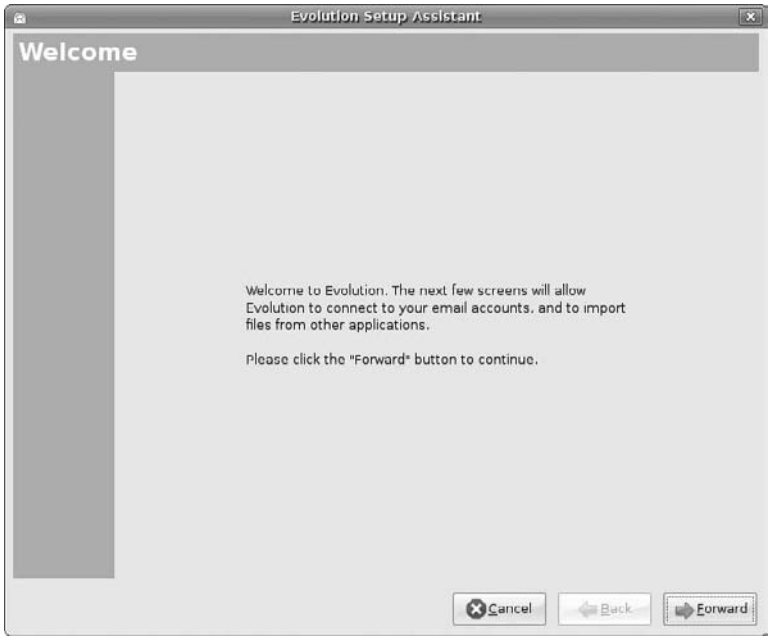
The next section discusses the Evolution Setup Assistant, which walks you through the process of creating your first e-mail account in Evolution.

Using the Evolution Setup Assistant

When Evolution actually starts up for the first time, the first window of the Evolution Setup Assistant, like the one shown in Figure 9-2, displays and walks you through the process of creating your first e-mail account in Evolution.

FIGURE 9-2

The Setup Assistant when starting Evolution for the first time



A slight clarification here, just to be on the safe side—the Evolution Setup Assistant configures Evolution to send and receive e-mail for an existing e-mail account somewhere; it doesn't actually create the account. For example, even though I use my Ubuntu machines to send and receive mail, my e-mail account is actually located at the ISP that hosts my domain. The Evolution Setup Assistant lets me tell Evolution about where it should look to retrieve my incoming e-mail and the mail server through which it should send outgoing mail. In my case, the machine on which I'm running Evolution is separate from the machine through which I actually send and receive my e-mail.

NOTE **Of course, if you run your own mail server and it's running on the same machine where you're running Evolution, your Ubuntu account may well be your official e-mail account. I discuss this in more detail in Bonus Chapter 9, "Setting Up a Mail Server."**

The idea of receiving e-mail can be confusing, so I generally think of this as the difference between *receiving* and *retrieving*—my actual e-mail account (at my ISP) receives my e-mail, and Evolution lets me retrieve my e-mail from there in order to read it, archive it locally or delete it, and so on. In official mailer-speak, this is the difference between a Mail Transfer Agent (MTA), a Mail Delivery Agent (MDA), and a Mail User Agent (MUA). The MDA at your ISP takes mail received by the MTA and puts it in your mailbox at your ISP. An MUA (in this case Evolution) retrieves a piece of mail from your mailbox at your ISP and displays it on your local system. (For more detailed information about these terms, see Bonus Chapter 9.)

Once you see the window in Figure 9-2, click Forward to move to the next portion of the setup process. A window like the one shown in Figure 9-3 displays.

FIGURE 9-3

Restoring from backups



If you are already an Evolution user and have created an Evolution backup file on another system, you can import the contents of that file by clicking the Restore Evolution from backup file checkbox and navigating to the location of your Evolution backup file using the navigation dialog that is displayed after clicking the Please select an Evolution Archive to restore button. You can create these backup files in recent versions of Evolution by selecting the File menu's Backup Settings command and specifying the name of the backup files that you want to create. These backup files are created as compressed tar files, with a default name of `evolution-backup.tar.gz`.

If you are already an Evolution user and have an archive file, you can skip ahead to the "Sending and Receiving Mail" section of this chapter after restoring the backup file. No further configuration is necessary.

Migrating from Older Versions of Evolution

If you are already an Evolution user but are using a version of Evolution that does not provide the Backup Settings command, you can still migrate your existing Evolution configuration information by using the tar command to create a compressed tar file that contains the contents of the .evolution directory in your home directory (and its subdirectories). You could do this by using commands such as the following:

```
cd ~ ; tar czf evolution-backup.tar.gz
```

You can then transfer the file evolution-backup.tar.gz to your new Ubuntu system and restore it during the Evolution setup process. Do not delete your existing evolution information until you've verified that the backup file was successfully imported!

Older versions of Evolution used a directory called evolution in your home directory rather than the .evolution directory. If you're using version of Evolution that uses the evolution directory, congratulations for being a long-time Evolution user, but your upgrade mechanism will be a bit more complex. You'll have to back up this directory to a file, copy that file to your new Ubuntu system, and extract its contents manually before beginning the Evolution setup process. If you haven't done this, exit the Evolution setup assistant now and do so. When you restart Evolution's setup process after restoring your old evolution directory, the Setup Assistant should detect this old directory and offer to migrate everything to the new `.evolution` directory.

If you are new to Ubuntu and Evolution and therefore don't have anything to restore, click Forward to proceed to the next dialog. A screen like the one shown in Figure 9-4 displays.

The Full Name field should already be filled out for you—Evolution grabs this information from the account from which you're running Evolution. You can change that name, of course, if you'd prefer to have the people who receive mail from you see a name that's more "clever" than your actual account name. The critical bit of information on this panel is in the Email Address field. This must be the e-mail address that you want to associate with this account. As mentioned previously, this is typically your e-mail address at your ISP or wherever else you get your e-mail through.

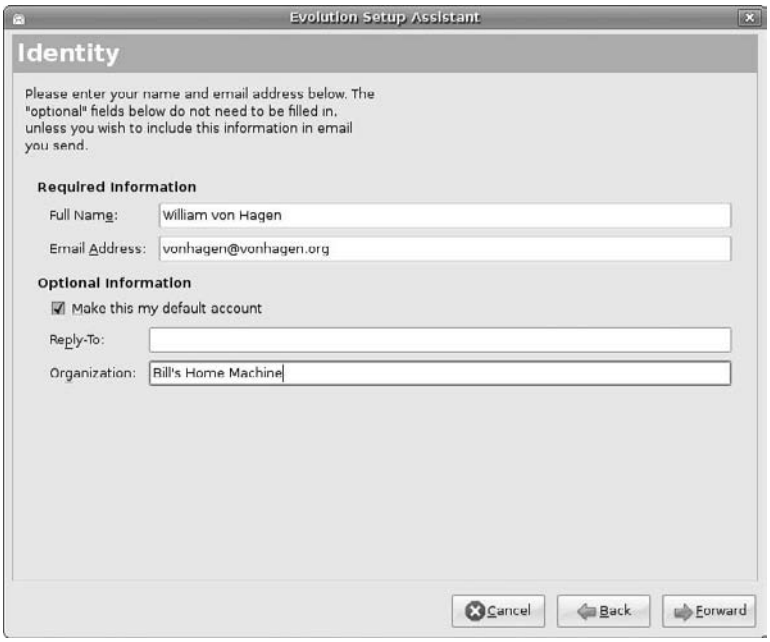
The Make this my default e-mail account checkbox is selected by default in this dialog because Evolution starts the Setup Assistant only when it detects that you are starting Evolution for the first time. Your default account is simply the one that Evolution will continue to use unless you activate another account.

If you want people who receive e-mail from this account to actually reply to another, you can specify the e-mail address of that account in the Reply-To field, but this is usually unnecessary. Similarly, if you want mail headers to announce that your mail is coming from a specific organization, you can specify the name of that organization in the Organization field. This is invisible to most recipients,

but it can be handy for filtering e-mail messages (explained later in this chapter) if you use this field to uniquely identify that this mail was coming from a specific location. I typically put something like "Bill's Home Machine" in this field, so that I can tell that I actually sent a specific piece of mail from my machine at home, if necessary.

FIGURE 9-4

Specifying your name and primary e-mail address



Once you've filled out this panel, click Forward to display the next panel in the Setup Assistant, shown in Figure 9-5.

When you first see this panel, it will be largely empty and the Server Type field displays the type None. Click the Description field to display a list of the different types of mail servers that Evolution can talk to, and select the one that corresponds to the primary e-mail protocol supported by the mail server at the location that actually hosts your e-mail, which is typically your ISP. Most mail servers nowadays speak POP (Post Office Protocol) or a special version of POP known as POP3, which adds some extensions to the standard POP mail server capabilities. The second most popular protocol nowadays is IMAP, which also includes Microsoft Exchange mailers, because they speak a variant (of course) of IMAP.

FIGURE 9-5

Specifying incoming mail server (MTA) parameters



A Quick Guide to POP vs. IMAP

The key difference between the POP and IMAP e-mail protocols is in where your e-mail is actually stored. POP mail clients typically download your e-mail from the POP server and store it on the machine where your mail client is running. This is extremely convenient for the operators of the mail server—they don't have to provide huge amounts of storage for e-mail because the storage requirements for POP mailers are transient. As people read their mail, it is downloaded to the client and then usually deleted from the server. IMAP servers, on the other hand, permanently store your mail on the server and only download the mail headers to your mail client when you first check your mail. These headers include the subject line, information about the sender, the date on which the message was sent, and so on. When you select a message to be read from an IMAP server, the mail is copied to your local system, but the master copy of the message remains on the IMAP server. The storage requirements for IMAP mailers therefore tend to increase over time, as more and more people accumulate larger collections of saved e-mail.

continued

continued

Each of these approaches has advantages and disadvantages. POP lets you keep all of your mail on your local machine, which is handy if you want to review your existing mail when you're not connected to the Net. The downside, of course, is that all of your mail lives on a single machine and it's therefore tricky to read old mail from a machine other than the one where that mail is actually stored. IMAP mailers make it easy to access your old mail from any machine, anywhere, but require a network connection to that machine in order to do so. Most IMAP clients enable you to synchronize your local mailboxes with those stored on the IMAP server, so that you can easily access old mail while traveling without requiring a network connection. However, ISPs that use IMAP hate people like me, who have huge mail archives going back ten years or so, and therefore require significant storage on an IMAP server.

Most POP mail clients, including Evolution's POP support, provide a special option called Leave mail on server. This enables you to configure multiple clients to read your POP mail, but only one of them (typically your home machine) actually downloads the mail and deletes it from the server. All POP clients remember the last message that they read, so they don't re-download all of your mail each time you check your mail—only those that they recognize as being new. This provides a nice compromise between the flexibility of IMAP and the local permanence of POP. For example, when I'm on the road, the copy of Evolution on my laptop lets me read my mail but doesn't delete it from the server—this only happens when I get back home and check mail from my primary machine. The downside of POP while traveling, of course, is that I don't have access to all of my old mail when I'm on the road, but if I'm really desperate, I can always connect to my home machine and either manually look through my mailboxes or run Evolution from that machine over the Net. Slow, but successful.

Once you select a specific type of server, the panel changes to reflect the particular types of settings that are associated with a mail server of that type, as shown in Figure 9-6, which shows the setting for a POP mailer.

On this panel, enter the name of your mail server and your username as far as the mailer is concerned. Note that your mailer's idea of your username may not actually be the same as your login name there—many POP mailers require a username of the form *login%domain.name* to guarantee the uniqueness of usernames on mailers that support multiple domains. For example, my login at my ISP is *vonhagen*, but my POP username is *vonhagen%vonhagen.org*.

Next, specify the security settings required to log in on your mailer—whether or not to try to use a secure connection to retrieve your mail from the mail server (which I always set to Whenever Possible), and the type of authentication required for connection to the server. Evolution provides a convenient Check for Supported Types button that you can click to probe your mail server to see what type(s) of authentication it supports. In most cases, you'll use Password authentication, but this option makes it easy to check whether your mailer supports more sophisticated types of authentication such as Kerberos or GSSAPI.

You can select the Remember password checkbox to optionally specify that Evolution remember the password to your incoming mail server, but I don't recommend that because this means that anyone who has access to your account will be able to read your mail.

Once you're finished specifying the options for receiving and retrieving incoming mail, click Forward to proceed to the next panel. A panel like the one shown in Figure 9-7 displays.

FIGURE 9-6

Specifying parameters for an incoming mail server

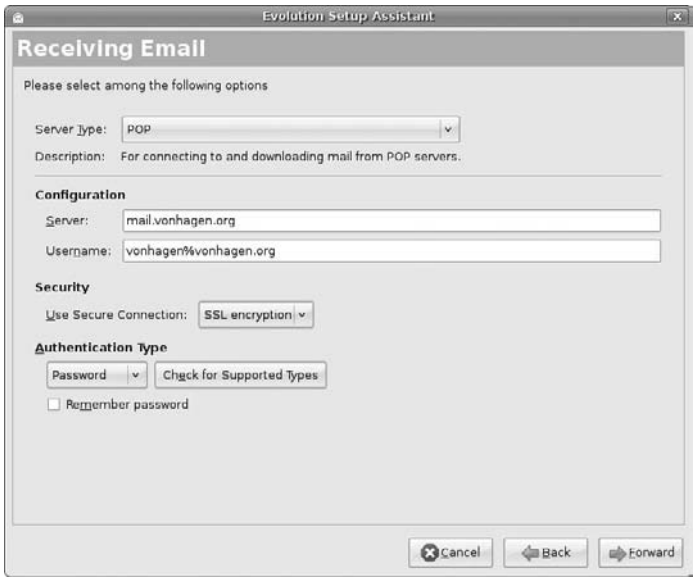


FIGURE 9-7

Options for handling incoming mail from a POP server

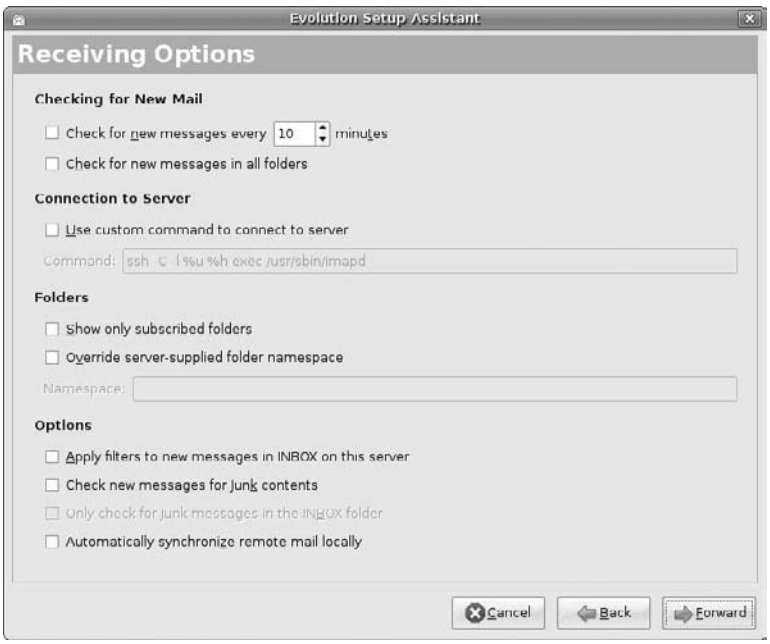


Panels like the one shown in Figure 9-7 enable you to configure the available options for how your mail client should interact with your incoming mail server. For all types of incoming mail servers, you can specify if, and how often, your mail client should automatically check for new e-mail. When configuring how Evolution interacts with an incoming POP mail server, you can also identify whether the copy of Evolution that you're configuring should remove the mail. As discussed in the sidebar "A Quick Guide to POP vs. IMAP" earlier in this chapter, POP mailers typically provide short-term storage for mail, while other types of mailers, such as IMAP, provide long-term storage for your mail. This panel also lets you configure whether Evolution should pay attention to POP3 extensions, which are site-specific extensions that the administrator of the mail server may have implemented.

Many of the options for interacting with an incoming mail server are specific to the type of server that you're connecting to. For example, Figure 9-8 shows the configuration options available when communicating with an IMAP server.

FIGURE 9-8

Options for handling incoming mail from an IMAP server



There are many more options for configuring Evolution's interaction with an IMAP server than with a POP server. Because IMAP servers maintain a master repository for your e-mail, Evolution's IMAP configuration panel offers options for how to connect to the server (such as through a custom command that provides a secure tunnel), how to interact with the mail folders on the IMAP server (both your personal folder and public ones to which you may have access), whether to synchronize your local mail folder with those on the server (giving you a local copy of all of your mail), and whether mail filters and general junk mail (spam) processing should occur against your master mailboxes or just locally. I'll discuss mail filters

and junk mail handling in the section "Filtering Incoming Mail" later in this chapter. If you're configuring Evolution to interact with an IMAP server, I suggest activating both the Apply filters…and Check new messages…options for now. You can always turn them off later, and selecting these options initially will ensure that Mail Filters work as expected, and will also help get rid of some of the spam that plagues every e-mail user nowadays.

After you configure how Evolution interacts with your incoming mail server when you read your mail, regardless of the type of incoming mail server you're using, you need to configure how Evolution actually sends mail. Click Forward to proceed to the setup panel shown in Figure 9-9.

FIGURE 9-9

Options for sending mail



Previous configuration steps helped you configure how your system receives mail and the characteristics of your incoming mail server. The panel shown in Figure 9-9 enables you to specify the characteristics of the mail server you use to send outgoing mail. Your incoming and outgoing mail servers may be different machines or processes because the tasks involved with sending and receiving mail are very different.

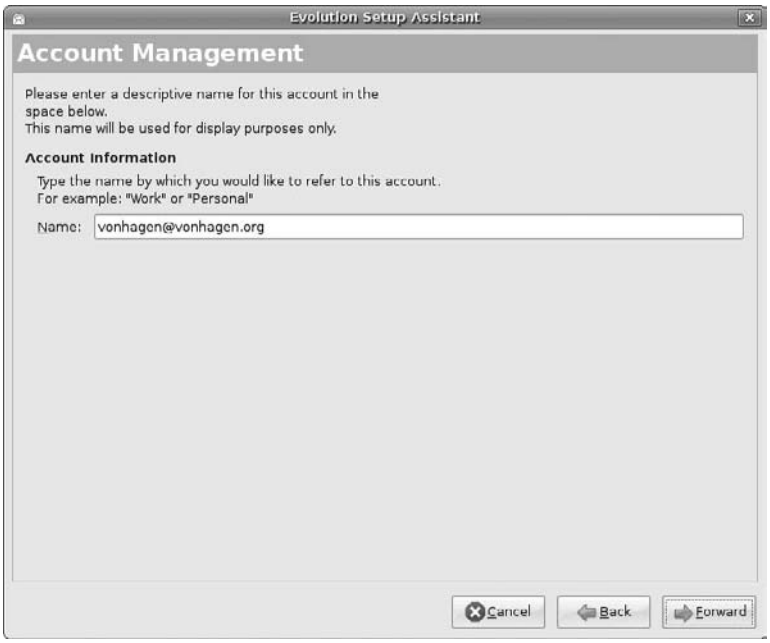
The critical piece of information on this panel is the type of your incoming mail server. Evolution supports two types of outgoing mail servers—SMTP (Simple Mail Transport Protocol), which simply means any mail server that conforms to the SMTP standard, and Sendmail, which is the most popular mail server used on Linux and other UNIX-like systems and runs locally on most Ubuntu systems. Most people will want to select SMTP—if you are running your own mail server on your system (and plan to continue to do so), you can select Sendmail, but if you are using an ISP for mail handling, you'll want to select SMTP. (For more information about running your own mail server, see Bonus Chapter 9, "Setting Up a Mail Server.")

If you select Sendmail, no additional configuration is required on this panel (because the server is running on the same machine as Evolution) but if you select SMTP, you'll need to specify the name of your outgoing mail server and provide information about how you authenticate to that server, if necessary. Most modern SMTP servers require authentication, so you need to check the Server requires authentication checkbox, specify that Evolution should use a secure connection whenever possible, and verify the type of authentication that your mail server supports by clicking the Check for Supported Types button, just as you did when providing similar information about your incoming mail server. As with the configuration of your incoming mail server, you can optionally specify that Evolution remember your password to the outgoing mail server, but I don't recommend that because this means that anyone who has access to your account will be able to send mail as you.

Once you've correctly configured Evolution's information about your outgoing mail server, click Forward to continue with the configuration process. Don't worry—you can always change these settings later if you discover that mail isn't working correctly, as explained in the next section. The panel shown in Figure 9-10 displays.

FIGURE 9-10

Naming a configured e-mail account



Almost there! The panel shown in Figure 9-10 simply requests a name for the e-mail account that you're configuring. This is the name that Evolution uses internally to identify this particular collection of e-mail settings, and is the name that you'll select in Evolution's Preferences dialog (discussed in the next section)

if you need to update or modify the settings for this e-mail account. By default, the name of a particular collection of settings is the same as the e-mail address for that account, which is what I typically use.

After making any desired changes to the name of this e-mail account, click Forward to continue. The panel shown in Figure 9-11 displays.

FIGURE 9-11

Specifying your time zone

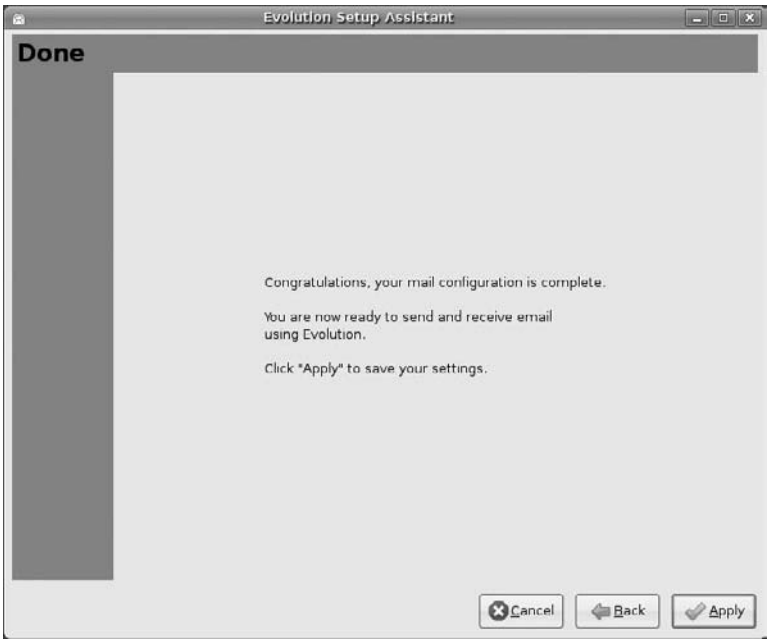


This panel enables you to specify the time zone that Evolution believes that you're running in, and thus uses to set the timestamps on your outgoing mail. By default, this panel displays UTC (Universal Time Coordinated), which means that Evolution believes that your system's clock is tracking time in location-independent, universal mode. Most people will want to specify their actual local time zone here, which you do by clicking on a location that is in the time zone as yours. You do this by moving the cursor over the map and clicking on a nearby location. As you move the cursor over the map, the status line at the bottom of this panel identifies the nearest city on the map that corresponds to the current cursor position. Clicking on any part of the map causes Evolution to zoom in on the associated portion of the map.

The Timezone map may not include the city that you actually live in. For example, I live near Pittsburgh, Pennsylvania, which (for some odd reason) isn't considered a major city, so I typically choose the nearest "real" city, which is New York. Once you've selected the nearest city to your actual location (or any city in the same time zone as yours), click Forward to proceed. The panel shown in Figure 9-12 displays.

FIGURE 9-12

Congratulations!



As you might gather from the panel shown in Figure 9-12, setting your time zone concludes setting up your initial mail account in Evolution. To save your settings, click Apply. Once you apply the current settings, you're ready to check your mail for the first time and begin using Evolution in general, as explained in the remainder of this chapter.

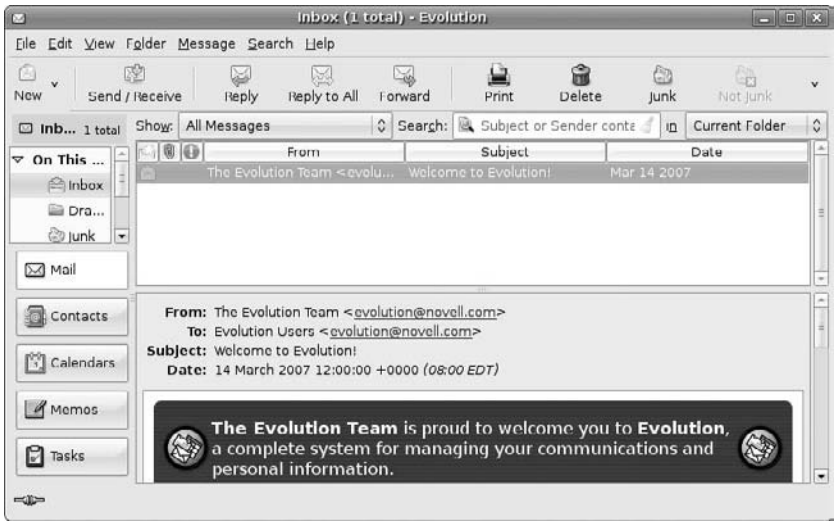
TIP **Unlike many e-mail clients, Evolution enables you to simultaneously send and receive mail for multiple e-mail accounts. Each e-mail account has a specific e-mail address, an associated server (and protocol) from which to read mail, and another through which to send outgoing mail. You can define other e-mail accounts using Evolution's Preferences dialog, which you can display by selecting the Preferences command from the Edit menu. One issue with Evolution's support for multiple POP e-mail accounts is that all addresses for which you retrieve and store mail in Evolution share the same set of local folders—each individual mail address doesn't have its own Inbox and set of folders. However, as I'll discuss later in this chapter, you can easily use mail filters or virtual folders to provide the appearance of having a per-address Inbox, which is the next best thing to having one.**

Sending and Receiving Mail

Figure 9-13 shows the window that you'll see when starting Evolution (once you've gone through its Setup Wizard to create an initial account). Let's read mail!

FIGURE 9-13

The primary Evolution window



To read your mail, click Send/Receive in the Evolution toolbar. Evolution prompts you for your password, contacts the remote mail server, and begins retrieving your mail. As it retrieves your mail (and sends any outgoing mail), Evolution displays a progress dialog, as shown in Figure 9-14.

FIGURE 9-14

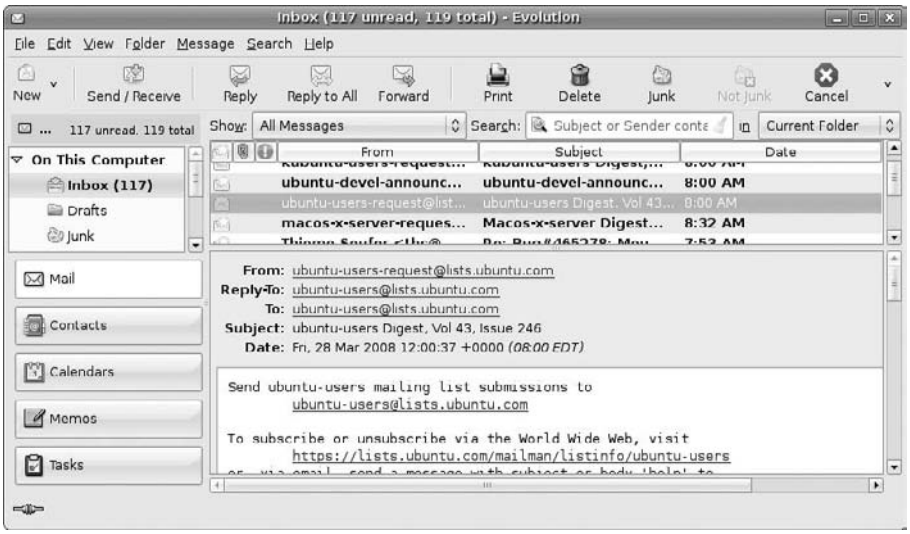
Evolution's e-mail progress bar



Once Evolution retrieves your incoming e-mail and sends any queued outbound mail, the progress dialog closes and Evolution updates its internal information about the state of the various mail folders on your system. To read your new mail, select your Inbox from the navigation pane at the left, and select any of the new mail messages to display that message, as shown in Figure 9-15.

FIGURE 9-15

Displaying a new message



Any folder shown in bold in the pane at the left is a folder that contains a message that is currently identified as being unread. Notice that the Evolution title bar displays the number of unread messages in your Inbox, as does the number in parentheses beside the Inbox folder in the navigation pane at the left.

The buttons in the Evolution toolbar are your key to the most common functions in Evolution:

- **New:** Creates a new item of any of the types of items supported by Evolution: mail messages, contacts, calendar items, memos, and tasks.
- **Send/Receive:** Retrieves any new mail held at your incoming mail server and sends any pending mail that you've sent.
- **Reply:** Replies to the sender of the currently selected e-mail message.
- **Reply to All:** Replies to the sender of the current message and sends a "carbon" copy (CC) of the mail message to anyone else who received a copy of that sender's message.
- **Forward:** Lets you forward the current e-mail message to a new recipient.
- **Print:** Opens a dialog that enables you to print the current message on any printer that has been configured on your system. (See Chapter 24, "Adding Hardware and Attaching Peripherals," for information on configuring printers.)
- **Delete:** Marks the current message for deletion, which removes the message from the list of active messages. To actually delete messages marked for deletion, select the File menu's Empty Trash command.

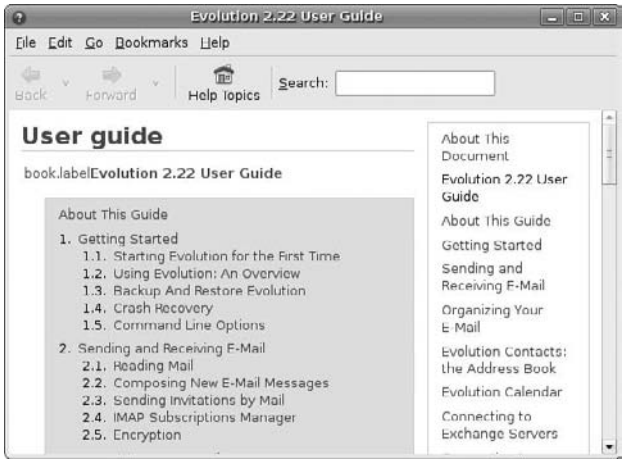
- **Junk:** Marks the current message as junk, reducing the chance that you'll receive similar messages in the future. For more information about junk mail handling, see the section "Automatically Checking for Junk Mail," later in this chapter.
- **Not Junk:** Enables you to identify a message that you had previously identified as junk as being acceptable. For more information about junk mail handling, see the section "Automatically Checking for Junk Mail," later in this chapter.
- **Cancel**: Cancels the current operation.

The down arrow at the end of the Evolution toolbar indicates that other tool buttons are available, but that the Evolution window on my screen currently isn't wide enough to display them all. You can click on the down arrow to select these other commands (Previous, and Next, by default) or simply resize the window to see these other tool buttons.

The buttons on the Evolution toolbar provide easy access to most of the commands that you'll ever need to work with your mail. However, Evolution is a large and powerful application with many bells, whistles, and knobs that you can turn, and deserves its own book. Luckily, Evolution also includes a great online manual that explains the ins and outs of using all of its capabilities. To access this online help, select Contents from the Evolution Help menu. Evolution's online manual displays, as shown in Figure 9-16.

FIGURE 9-16

Evolution's online help



Because this is a book on Ubuntu, not Evolution, the focus of this chapter is on helping you get started with Evolution. Until a book dedicated to Evolution is published, Evolution's online help and various Evolution-related Web sites are the ultimate source of information for this e-mail client. The next few sections provide some of my favorite tips for efficiently working with and managing your e-mail in Evolution. For additional information, see the Web sites listed in "Additional Sources of Information About Evolution" later in this chapter.

Sending Mail in Evolution

You can begin composing a new message in Evolution in a variety of different ways, including:

- Clicking the New button when you are viewing e-mail
- Pressing Ctrl+Shift+M on your keyboard
- Selecting Compose new message from the Message menu

You can reply to a message that you are currently viewing by doing either of the following:

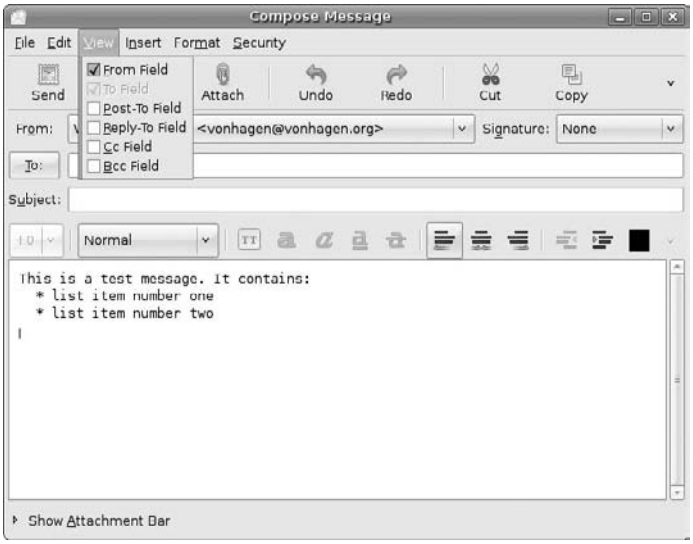
- Clicking the Reply button, pressing Ctrl+R on your keyboard, or selecting the Reply command from the Message menu.
- Clicking the Reply to All button. Pressing Shift+Ctrl+R on your keyboard, or selecting the Reply to All command from the Message menu. This has the same effect as Reply if you are the sole recipient of the currently selected message.

TIP **If you are currently viewing a message that was sent to a mailing list to which you are a subscriber, you can press Ctrl+L or select the Message menu's Reply to List command to reply to the list rather than to the individual who posted that specific message. To follow the rules of good Netiquette, always trim as much of the message you're replying to as possible—especially if you're replying to a mailing list that you receive in digest form.**

Selecting any of these options displays Evolution's Compose Message window, as shown in Figure 9-17. If you are replying to an existing message, the contents of that message are displayed in the composition window in quoted form to provide some context for your response.

FIGURE 9-17

Composing a message in Evolution



When you're composing e-mail, Evolution initially displays only the From and To fields in the message header to save space. In many cases, you'll want to send a copy of your mail to someone else (or perhaps to another e-mail address of yours). To display the fields of the mail header that support such additional recipients, click the View menu and select the other fields that you want to see such as CC (Carbon Copy), BCC (Blind Carbon Copy—one or more e-mail addresses that aren't listed in the mail that is actually sent to the other recipients), and so on.

Once you've finished composing your e-mail message, click Send to send the message. This queues the message to be sent, and you can continue working with your e-mail—while you're working, Evolution contacts the remote mail server in the background and sends the message as quickly as it can.

Setting Evolution Preferences for Sending and Viewing Mail

Evolution's Preferences dialog enables you to configure Evolution, and also to correct or change any of the basic information that you supplied when initially configuring Evolution. This section discusses one of the most common configuration changes, which is the format in which you send your e-mail.

Two basic types of encodings are used when composing and receiving e-mail. One is plain text, which means that your messages display as standard text that might come out of a typewriter (if you remember what those are), with no special fonts or special formatting. The other is e-mail that is formatted as HTML, which lets you use different fonts, special characters for bullets, and so on. Many people prefer one or the other, and whether or not you can read mail delivered in HTML format is up to the e-mail client that you're using (although most can, nowadays).

NOTE **Whether or not an e-mail client *can* send mail in HTML format is a different question than whether you *should* send e-mail formatted as HTML. Many people, and many specific recipients such as mailing lists, do not want to receive or will not accept e-mail formatted as HTML. I'm not a fan, either—I always send text mail.**

As shown in Figure 9-17, when you first compose messages in Evolution, it uses standard text format for your messages by default. If you want to switch to using HTML mail, you can do this in either of two ways:

- If you want to send just the current message in HTML format, select the HTML command from the Format menu in the window where you're composing your message. Figure 9-18 shows the same mail composition window as that in Figure 9-17, but displayed as HTML. Additional formatting buttons are active when sending an HTML message, enabling you to use different fonts, different justification formats for the paragraphs in your messages, and so on.
- If you want to send all of your mail as HTML, you must set that as your preference in Evolution's Preferences dialog.

Evolution provides two Preferences panels for setting how it sends and interacts with incoming mail that is formatted as HTML. These are located on the Preferences dialog's Mail Preferences and Composer Preferences panels. To display Evolution's Preferences dialog, select the Edit menu's Preferences command.

To configure Evolution so that it always sends HTML mail, click Composer Preferences in the pane at the left. The panel shown in Figure 9-19 displays.

To always send HTML mail, click the checkbox beside Format messages in HTML. When you close Evolution's Preferences dialog, this will become the default format for all new messages that you create.

To set the options for how Evolution handles HTML-formatted mail that you receive, click Mail Preferences in the pane at the left and click the HTML Mail tab. The panel shown in Figure 9-20 displays.

FIGURE 9-18

A message formatted as HTML

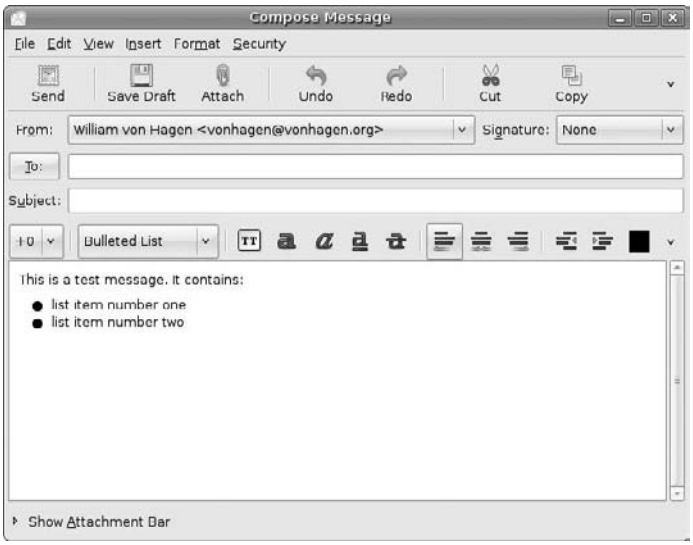


FIGURE 9-19

Specifying HTML as your default mail format

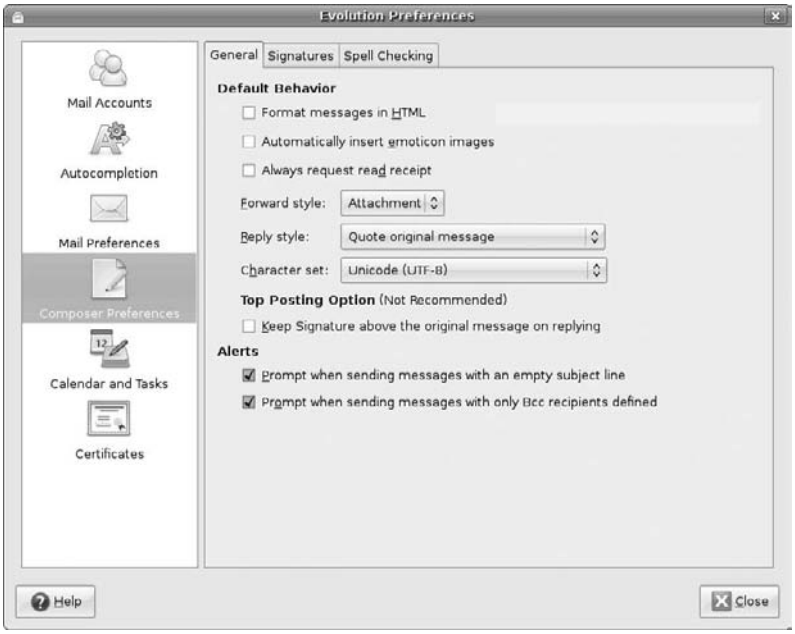
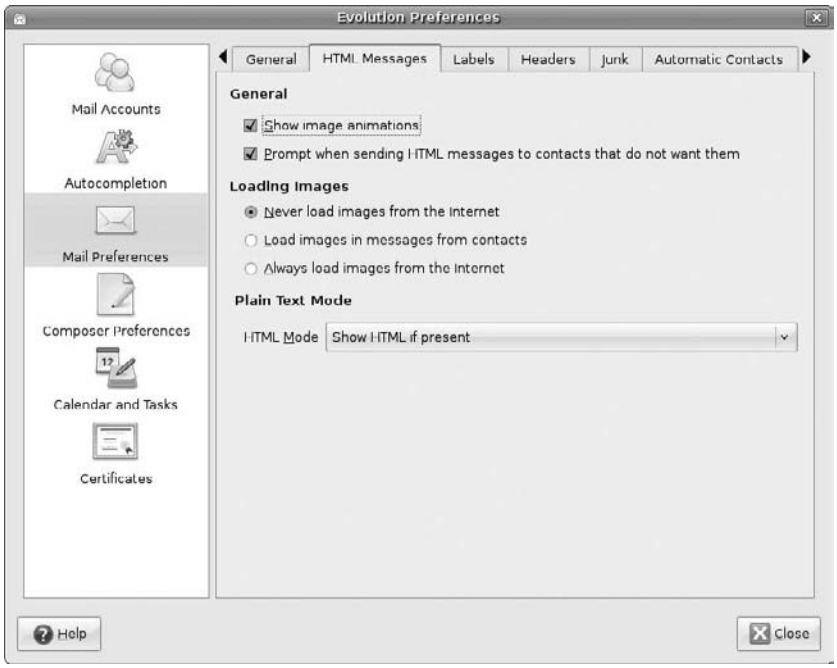


FIGURE 9-20

Options for displaying incoming HTML mail



This panel provides several options for incoming HTML mail. For example, if you are occasionally checking your personal e-mail on a machine other than your primary home machine, you may want that to happen as fast as possible. In that case, you do not want to load images, because reading mail that contains images is slower than plain mail because images take some time (and network bandwidth) to download. To stop Evolution from downloading images at all, select the Never load images from the Internet radio button. If you want to download images only from people that you know (i.e., people who are in your contacts list), select the Load images in mail from contacts radio button. If you always want to load images, select the Always load images from the Internet radio button.

The panel shown in Figure 9-20 also provides a nice option for the benefit of those entities that do not want to receive HTML e-mail: the "Prompt when sending HTML messages to contacts that do not want them" option. This causes Evolution to display a confirmation dialog if you are sending HTML mail to a person or mailing list that identifies itself as not wanting to receive HTML e-mail.

Once you've finished setting Evolution's mail preferences, click Close to exit the Preferences dialog. Most of the preferences that you can set take effect immediately. Some may not take effect until the next time you start Evolution.

TIP **Evolution enables you to customize many other aspects of how it composes new messages, displays incoming e-mail, interacts with remote mail servers, and many more. Explaining all of these is outside the scope of this chapter, but I encourage you to explore the various panels and tabs available in Evolution's Preferences dialog to make sure that Evolution works exactly the way that you want it to.**

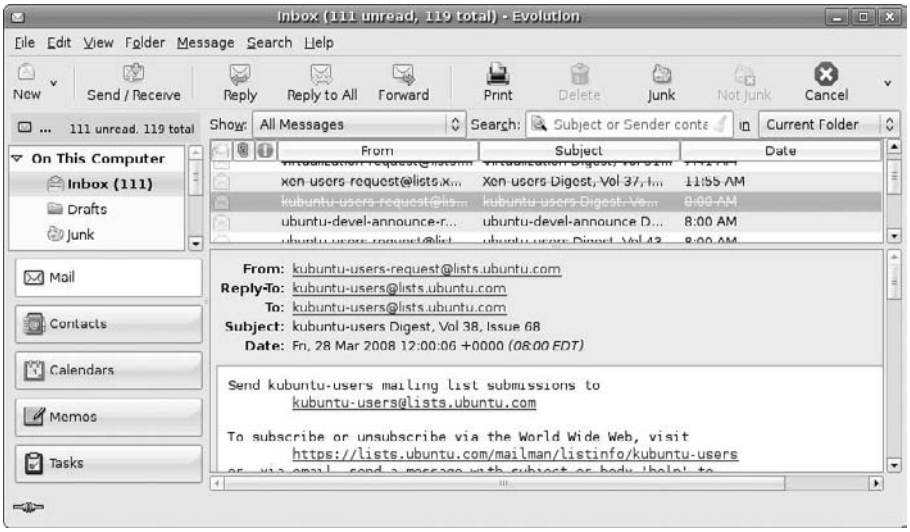
Undeleting Mail Messages

Like files deleted on DOS and Microsoft Windows systems, e-mail messages that you delete aren't immediately purged, but are just marked for deletion when you explicitly empty the Trash or exit Evolution (if you've configured Evolution to do that).

If you accidentally delete an e-mail message and would like to get it back, you can deselect the View menu's Hide Deleted Messages option. After selecting this option, the titles for messages that are marked for deletion are shown with a line through them, as shown in Figure 9-21.

FIGURE 9-21

Showing mail messages that are marked for deletion



You can undelete any mail messages that are currently marked for deletion by selecting the message title and selecting the Edit menu's Undelete message command.

Creating and Using Mail Folders

After sending and receiving a fair amount of e-mail, your Inbox will become pretty cluttered and you'll find it hard to locate a specific message even though Evolution provides some nice options for locating specific messages above the window where your Inbox is displayed. But why search for something if you can organize things better in the first place? The list of mail folders that Evolution provides by default is always displayed in the navigation pane at the left of your mail Evolution window, but Evolution also makes it easy for you to create your own custom folders to better (and permanently) organize your mail.

Evolution creates the following folders for you during its configuration process:

- **Inbox:** The folder in which incoming mail is deposited when read from your mail server.
- **Drafts:** A folder in which you can put mail messages that you're working on but which have not yet been sent.

- **Junk:** A folder that holds messages that have been identified as junk messages by Evolution's junk-mail scanner. For more information about this folder and junk mail in general, see the section "Automatically Checking for Junk Mail," later in this chapter.
- **Outbox:** A temporary location where mail that you have sent is stored before it is actually sent to your outgoing mail server.
- **Sent:** A folder that is provided to hold copies of outgoing mail that you have sent, so that you have a record of their contents. Copies of sent messages will only be saved if you configure Evolution to do so using the Defaults tab in the Edit ➪ Preferences dialog's Account Editor.
- **Trash:** A folder that contains messages slated for deletion when you explicitly empty the Trash or exit from Evolution (if you've configured your Evolution preferences to perform that action).

With the exception of your Inbox, mail folders are listed alphabetically. Your Inbox is always the top folder in Evolution's navigation pane, so that it is always easy to find.

You can create a new mail folder by doing either of the following:

- Selecting New from the Folder menu
- Right-clicking the folder or location in which you want to create the new folder and selecting New Folder from the pop-up menu

After performing either of these actions, a dialog like the one shown in Figure 9-22 displays.

FIGURE 9-22

Creating a new mail folder



In this dialog, you can navigate to the location where you want to create the mail folder, enter the name of the new folder, and click Create.

Creating special-purpose mail folders makes it easy to organize your mail so that your Inbox contains only a minimum number of messages—often, only new ones. Mail folders can be hierarchical—in other words, the folder can also contain other folders. For example, I typically create a Friends folder with subfolders that correspond to each of the people with whom I regularly communicate. Once I've read or responded to a message, I put it in the appropriate folder. This makes it easy to review my correspondence with specific people without having to sift through many unrelated messages.

Using Search Folders

In addition to standard mail folders for organizing your mail, Evolution was the first mail client to introduce the idea of search folders for viewing related groups of mail messages. Standard mail folders actually hold mail messages—search folders are like an index or table of contents to a group of related messages that contains pointers to those message rather than the messages themselves. This makes it easy to store your messages in explicit mail folders but use virtual folders to give you alternate views of your mail folders. Search folders were known as *vfolders* or *virtual folders* in older releases of Evolution.

As an example, suppose that you subscribed to various mailing lists about different Linux distributions and types of Linux filesystems, and organized your incoming mail based on the name of the mailing list. Suppose that you also wanted to be able to examine all of the messages about a specific type of filesystem, regardless of which list they were posted on. You could always create a mail filter (as explained in the next section) to file the messages according to the mailing lists that they arrived on and copy them to a filesystem-specific folder, but that would eventually end up wasting a significant amount of space. Evolution's search folder concept to the rescue—it is easy to create a search folder that simply contains pointers to all of the matching messages without duplicating them.

Like most tasks in Evolution, there are a variety of ways to create a new search folder. These include the following:

- Right-click on the Search Folders item in the navigation pane and select the New Folder command from the pop-up menu.
- Right-click on the heading on any sample message that satisfies the criteria for the search folder that you want to create, select the Create Rule From Message submenu, and select any of Search Folder from Subject, Search Folder from Sender, Search Folder from Recipient, or Search Folder from Mailing List. (The last displays only if you select a message that Evolution can recognize as having been sent to a mailing list.) Each of these menu commands creates an initial search folder based on searching the specified portion of that mail message, which you can retain as is or further customize.
- Select the Edit menu's Search Folders command and click Add to create a new search folder.
- While any sample message that satisfies the criteria for the search folder that you want to create is highlighted, select the Message menu's Create Rule submenu, and select any of Search Folder from Subject, Search Folder from Sender, Search Folder from Recipients, or Search Folder from Mailing List. (The last only displays if you select a message that Evolution can recognize as having been sent to a mailing list.) Each of these menu commands creates an initial search folder based on searching the specified portion of that mail message, which you can retain as is or further customize.

Performing any of these actions displays a dialog like the one shown in Figure 9-23 that contains any initial criteria implied by the command that you selected.

Once this dialog displays, you can specify a name for the search folder and add or refine the selection criteria for the search folder in the same way that you create filter rules for incoming mail, as explained in the next section. The primary difference between search folder and filter rule creation lies in the bottom half of the dialog shown in Figures 9-23 and 9-24, which (for search folders) identifies the location(s) that the folder should search for matching messages. You can specify that only certain folders be searched by clicking Add and selecting their names, or you can select With all Local Folders to include messages from any folder, or With all Local and Remote Folders to search the server-side, remote folders supported by mail protocols such as IMAP. Figure 9-24 shows a search folder that has been defined to search only a specific set of local folders.

FIGURE 9-23

Creating a new search folder

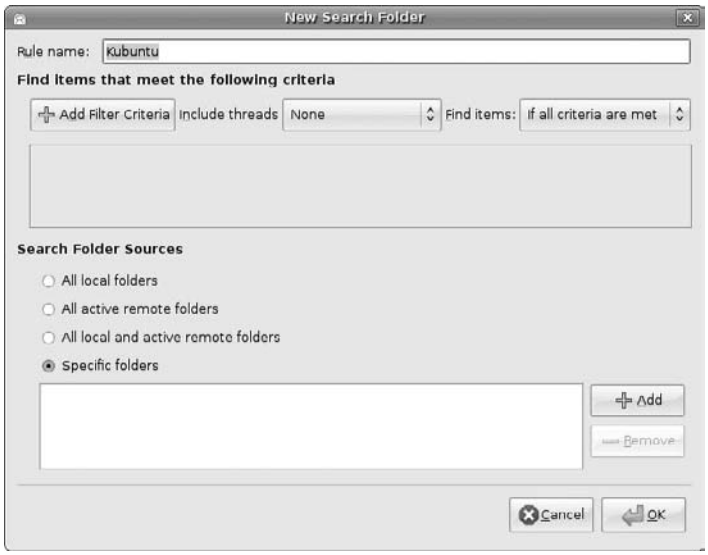
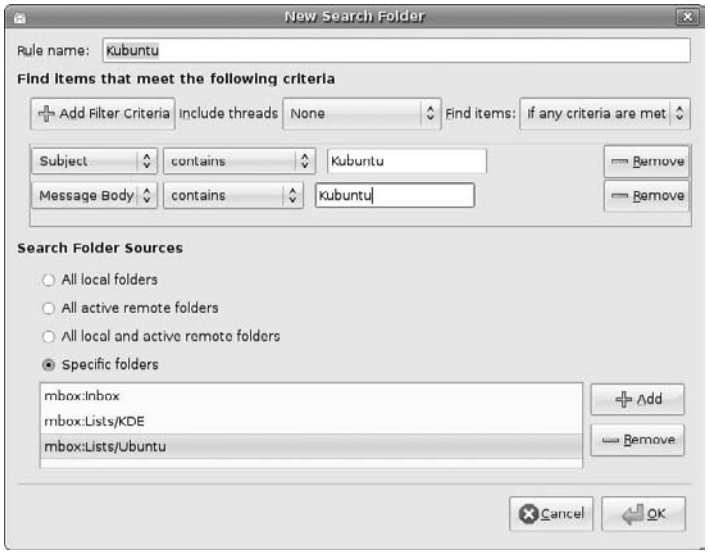


FIGURE 9-24

A search rule associated with specific local folders



Once you're finished fine-tuning the rule that messages must match to be included in your new search folder, click OK to close this dialog and return to Evolution. The search folder is created under the Search Folder item in Evolution's navigation pane, and contains pointers to any existing messages that match the specified criteria.

Filtering Incoming Mail

We all receive a significant amount of e-mail nowadays, some of which we actually want to receive. It's a sad fact of modern Internet life that all of us receive a significant amount of spam or junk mail. These are essentially equivalent terms for mail that we didn't ask for, but which is either broadcast to collections of randomly generated e-mail addresses or perhaps sent to us directly if we were unlucky enough to get our e-mail addresses included on one of the lists that some companies, and certainly spammers, sell to each other.

Evolution provides two sophisticated features for automatically analyzing incoming mail and performing specific actions on mail messages that match these criteria. In the case of legitimate mail, Evolution enables you to create powerful filtering rules that can automatically file incoming mail in appropriate folders. Mail filters can also discard incoming messages from known spam sites, messages that contain words that only a spammer would use, or messages that match any other criteria that you define as spam.

In addition to making it easy for you to create powerful filtering rules, Evolution can also be integrated with other free Linux software that uses different approaches to identify spam and junk mail and automatically move it to your Junk folder. You can then quickly scan your Junk folder occasionally to see if all the mail in there is indeed junk and, if so, you can easily delete it without having to read it.

The next two sections explain how to create your own filter rules and how to install the software required for junk mail processing in Evolution on Ubuntu systems.

Automatically Processing Incoming Mail

To expand on the definition of filtering given in the introduction to this section, filtering e-mail means to analyze its contents based on various criteria and to use the results of that examination to perform some action regarding that message. Message filters are typically applied to incoming messages that are stored locally—if you want to apply message filters to messages stored in remote folder (such as your Inbox on an IMAP server), you must specify that as an option when creating your account, as shown in Figure 9-8.

As with most tasks in Evolution, there are several ways to create a new mail filter. These include the following:

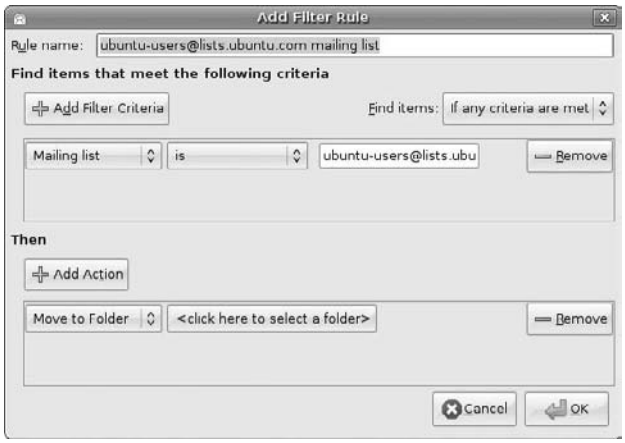
- Right-clicking on the heading on any sample message that satisfies the criteria for the message filter that you want to create, selecting the Create Rule From Message submenu, and selecting Filter on Subject, Filter on Sender, Filter on Recipients, or Filter on Mailing List. (The last displays only if you select a message that Evolution can recognize as having been sent to a mailing list.) Each of these menu commands creates an initial message filter based on the specified portion of that mail message, which you can retain as is or further customize.
- Selecting the Edit menu's Message Filters command and clicking Add to create a new message filter.
- While any sample message that satisfies the criteria for the search folder that you want to create is highlighted, selecting the Message menu's Create Rule submenu, and selecting any of Filter on

Subject, Filter on Sender, Filter on Recipients, or Filter on Mailing List. (The last displays only if you select a message that Evolution can recognize as having been sent to a mailing list.) Each of these menu commands creates an initial message filter based on the specified portion of that mail message, which you can retain as is or further customize.

Performing any of these actions displays a dialog like the one shown in Figure 9-25 that contains any initial criteria implied by the command that you selected.

FIGURE 9-25

Creating or editing a message filter



When creating a message filter, the default action to perform on matching messages is to move that message to a specified folder. To identify that folder, click <click here to select a folder> and locate the target folder in the navigation dialog that displays.

Evolution provides a very sophisticated set of possible ways of matching messages in a message filter. The following are the portions of the message that you can match and the comparison operators that you can use. These correspond to the contents of the drop-down menus you see when clicking on the first two buttons when creating or editing a rule:

- **You can examine the following portions of a mail message when creating a rule:** Attachments, Date Received, Date Sent, Expression, Follow Up, Junk Test, Label, Match All, Message Body, Mailing List, Pipe to Program, Recipients, Regex Match, Score, Sender, Size (KB), Source Account, Specific Header, Status, and Subject.
- **You can analyze those portions of an e-mail message using the following comparison operators:** Contains, Does Not Contain, Does Not End With, Does Not Sound Like, Does Not Start With, Ends With, Is, Is After, Is Before, Is Greater Than, Is Less Than, Is Not, Returns Greater Than, Starts With, Sounds Like, And Starts With. Some of these operators can only be used when examining specific portions of a mail message. As an example, the "Is After" and "Is Before" operations can only be used with dates.

When creating a rule for matching a filter, you can select any of the items in the first and second columns, in any combination. Not all of the items in the second column are available for all types of match conditions—for example, options such as "is greater than" and "is less than" apply only when you are checking a match condition, such as Size, that has a numeric value.

As an example of using match conditions for filter rules, to create a rule that matches when a recipient of an e-mail message was ubuntu-users@lists.ubuntu.com, you would select Recipients from the first drop-down menu in the dialog shown in Figure 9-25, select is from the second drop-down rule menu in that figure, and enter ubuntu-users@lists.ubuntu.com in the text area at the right to complete the rule.

TIP **When creating match rules, the most important control to set is whether a message must match only one (or more) of the specified criteria, or whether the associated operations will be performed only if the message matches all of the specified criteria. This can be the difference between a successful rule and one that doesn't work as you intended. This is controlled by the Find Items control, shown in the upper-right corner of any message filter dialog, which you can set to either If any criteria are met or If all criteria are met.**

Table 9-1 shows a list of the possible operations that you can perform on messages that match any (or all) of your filter conditions.

TABLE 9-1

Actions to Perform on Matching Messages

Operation	Action
Adjust Score	Select value from 3 to –3 or enter number to add.
Assign Color	Select color.
Assign Score	Select value from 3 to –3 or enter number.
Beep	N/A.
Copy to Folder	Select folder.
Delete	N/A.
Move to Folder	Select folder.
Pipe to Program	Select program.
Play Sound	Select sound file.
Run Program	Select program.
Set Status	Select from Replied to, Deleted, Draft, Important, Read, and Junk.
Stop Processing	N/A.
Unset Status	Select from Replied to, Deleted, Draft, Important, Read, and Junk.

In Table 9-1, the first column is the action that you want to perform, and the second column enables you to specify the value for that operation. For example, if you want to move a message to a specified folder after matching a filter rule, select the Move to Folder operation from the first drop-down menu in the Then portion of the rule, and click <click here to select a folder> to identify the target folder for that operation.

As an example of creating the Then portion of a filter rule, I'll continue with the example of messages whose recipients match ubuntu-users@lists.ubuntu.com. Assuming that I want to move such messages to their own folder, I would click Move to Folder from the drop-down menu in the first column and then navigate to the correct target folder. I could also add extra actions, such as Setting the status of the message to Read (otherwise the folder will show up in bold, indicating that it contains one or more unread messages), and conclude the rule with a Stop Processing instruction, which ensures that no other filter rules will be applied to the current message. This generally expedites message filtering, because once a rule matches a message, no other analysis of that message is performed and Evolution moves on to checking the next message.

Automatically Checking for Junk Mail

Evolution's ability to filter e-mail messages based on criteria that you can define is powerful, but requires that you create a potentially huge number of rules to catch all of the spam and other junk mail that you may receive. Luckily, Evolution also supports the integration of external applications that can analyze messages and identify them as junk if that's the case. Popular examples of this sort of junk e-mail filter application are SpamAssassin (`http://spamassassin.apache.org/`) and Bogofilter (`http://bogofilter.sourceforge.net/`).

Evolution includes a plug-in for SpamAssassin, which will automatically execute the SpamAssassin program and apply it to incoming mail. Because SpamAssassin has special licensing requirements, the Ubuntu folks do not include SpamAssassin as part of a default Evolution installation. To install SpamAssassin on your system, you must manually enable the officially unsupported Universe repository and use Synaptic Package Manager or apt-get to install SpamAssassin to enable Evolution's default Junk mail filtering. I describe enabling different repositories and installing custom software packages in Chapter 23, "Adding, Removing, and Updating Software."

Once you have downloaded and installed SpamAssassin, check Evolution's plug-in list to make sure that SpamAssassin integration is active. To do this, select Plugins from the Edit menu to display the Plugin Manager dialog, shown in Figure 9-26.

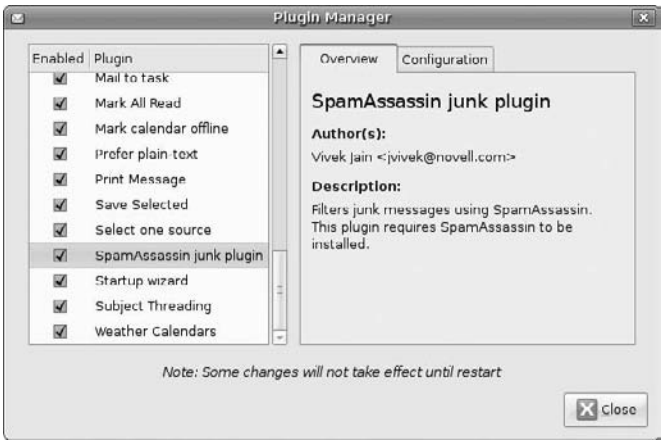
Once this dialog displays, scroll down to the SpamAssassin junk plug-in entry and make sure that it is enabled. Click OK after verifying this and enabling this plug-in, if necessary.

Once you install SpamAssassin and make sure that it is enabled, clicking the Junk button in the Evolution toolbar will identify the currently selected message as junk and will mail it to your Junk folder. Subsequent messages you receive that are conceptually similar to that message will automatically be identified as junk and also moved to the Junk folder. You should check the Junk folder periodically, verify that all of the messages that it contains are actually junk mail, and delete those messages. If you see messages classified as junk that are actually valid messages, select those messages and click the Not Junk button in the Evolution toolbar to make sure that similar messages are not categorized as junk. Over time, you will see that Evolution's support for junk mail filtering does a good job of identifying spam and other junk mail, saving you lots of time that would otherwise be spent examining such messages for criteria that you could use to manually filter out those e-mail messages.

TIP **If you are having problems installing or using SpamAssassin, make sure that Evolution's BogoFilter plug-in is not enabled on the Edit ➪ Plugins dialog. This conflicts with SpamAssassin and can cause problems.**

FIGURE 9-26

Enabling the SpamAssassin plug-in



Accessing Free and Web-Based E-mail

The broadband or dial-up connection that you use to connect to the Internet from your home is generally supported by an academic or commercial Internet service provider (ISP) that also provides you with an e-mail address. Such e-mail addresses are inherently transient—once you finish with a class, graduate, and can't keep your `.edu` address, or change commercial ISPs, your e-mail account there is usually deleted. This may cause some of your more high-strung friends and family members, used to regularly exchanging mail with you as bigdaddy@some-random-ISP.com, to panic and begin sending sympathy cards to your wife, when all the while you simply have a new e-mail address.

As an alternative, companies such as AOL, Google, Microsoft, and Yahoo!, as well as hundreds of specialized Web sites, provide free e-mail accounts that you can use regardless of your ISP, and which you can access via any Web browser. Of course, these accounts are basically just free advertising for these companies, driving traffic to their Web sites (`gmail.com`, `hotmail.com` or `live.com`, or `yahoo.com`, respectively), displaying advertisements in some cases, and so on.

With the exception of Gmail, you can read and send mail at these sites only by using a Web browser. This is a pain if you would like to use a single e-mail application, such as Evolution or Kontact, to read and manage all of your e-mail. However, Gmail offers both POP and IMAP access, and you can at least read and archive your Web-based e-mail from Evolution and Kontact (or any other mailer that supports POP/POP3 mail) by installing and configuring software such as FreePOPs. The next few sections explain each of these.

TIP **If you are interested in getting a free e-mail account, the EmailAddresses.com site (`www.emailaddresses.com/guide_types.htm`) is a good source of general information about the many different sites that provide free IMAP, POP3, and webmail accounts.**

Reading Gmail from a Desktop Mail Client

Google, providers of Gmail, is one of the few big-name free e-mail providers that actually support POP and IMAP. This section explains how to configure Gmail to support either POP or IMAP. You could, of course, activate both to give yourself IMAP access to your account until you drain your e-mail with a POP client at some time, but I've never found that to be necessary.

The basic steps for reading Gmail using Evolution (or any other mail client) are as follows:

1. Configure Gmail to support the mail transport protocol that you want to use.
2. Create an account in your mail client that gives you access to your Gmail mail.

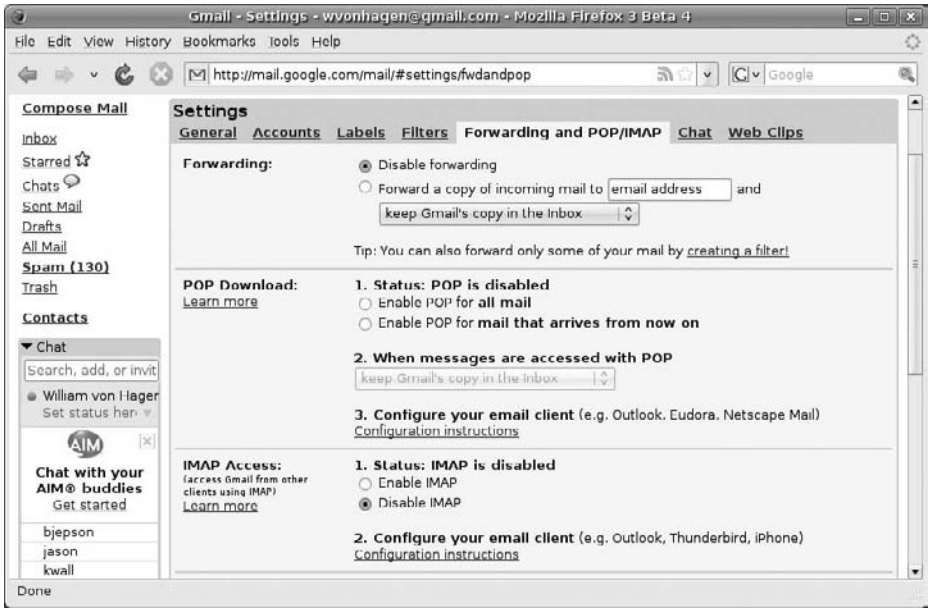
The next few sections provide the information that you'll need to do these steps for both POP and IMAP.

Selecting a Transport Protocol in Gmail

In order to access Gmail via POP or IMAP, you must first configure your Gmail account to be available via the protocol of your choice. To do this, log in to your Gmail account via Firefox—I mean, via your browser of choice—and select the Settings link in the upper-right portion of the screen. Next, select the Forwarding and POP/IMAP link in the header of the page that displays. This displays a screen like the one shown in Figure 9-27.

FIGURE 9-27

Configuring Gmail for POP/IMAP



The bottom two settings on this screen enable you to activate POP and IMAP connectivity to your Gmail account. Personally, I tend to favor IMAP with Gmail because this guarantees that all of my Gmail mail is archived at Google and enables me to access it from anywhere, including any folders that I've set up in my Gmail account. Because I tend to use a variety of systems, this gives me a single view of my Gmail mail from anywhere. Your mileage may vary.

TIP **If you decide that you want to read Gmail via the Web only, you can also activate Forwarding, which will send a copy of any incoming Gmail mail to another e-mail address. To do this, click the "Forward a copy of incoming e-mail to" radio button, enter your other e-mail address, and make sure that the second configuration list says to keep Gmail's copy in the Inbox. This provides a nice notification mechanism for Gmail, which can be handy if your Gmail account is not your primary e-mail account.**

To activate IMAP access to your Gmail account, simply click the Enable IMAP radio button and click Save Changes at the bottom of the screen.

If you decide to use POP rather than IMAP, a few more configuration choices are available. First, you can either enable POP for all of your mail, or for all of the mail that arrives after the point at which you configured POP access. Which of these you want to use depends on how you want to use Gmail:

- If you want to personally archive all of your Gmail correspondence on one of your own systems, you'll probably want to select the Enable POP for all mail. The first time that you access your Gmail account with a mail client that is not configured to leave your mail on the server, all of the mail in your Inbox will be downloaded to your desktop system, and what happens to it in your Gmail account depends on how you set the When messages are accessed with POP configuration option (explained later in this section).
- If you use this option in the hopes of archiving all of your Gmail mail on one of your local systems, note that POP only deals with mail in the Inbox. Any mail that you've archived in folders that you created in your Gmail account will not be downloaded, only mail in the Inbox. If your goal is to archive everything locally, you'll have to copy or move the messages from your folder into your Inbox. Copying them will leave a copy in your Gmail account, but you won't be able to access those folders via POP.
- If you are simply using your mail client to check your Gmail mail and occasionally send messages but have configured your mail client to leave messages on the server, the first time that you check mail on any new system, you'll see a deluge of all of the mail that is still in the Inbox for your Gmail account. In this case, you should probably select the "Enable POP for all e-mail that arrives from now on" radio button.

There are plenty of other scenarios in which you may prefer one option over the other—these are just the most common cases. Personally, I like IMAP.

If you're still sold on POP, you can use the When messages are accessed with POP configuration item to fine-tune what happens to your Gmail mail when you read it via POP. These settings are as follows:

- **Keep Gmail's copy in the Inbox:** The default setting, this value enables you to read your Gmail mail over the Web but still access it via POP. This setting overrides whether your mail client wants to delete messages that it has downloaded, and can protect you if you forget to tell your mail client to leave the messages on the server.
- **Delete Gmail's copy:** This value deletes mail from your Gmail Inbox once you've read it via POP, and overrides the settings in your mail client. In other words, even if you've configured your mail client to leave messages on the server, any Gmail mail that you access via POP will be purged from your Gmail account.

- **Archive Gmail's copy:** This value deletes mail from your Gmail Inbox once you've read it via POP but leaves an archive copy of the mail on your Gmail account that you can still access via the All Mail link. This setting also overrides the settings in your mail client. In other words, even if you've configured your mail client to leave messages on the server, any Gmail mail that you access via POP will be purged from your Inbox but will still be available in the All Mail link.

After you have activated POP or IMAP on this screen, optionally enabled Forwarding, and specified any related configuration parameters, click Save Changes at the bottom of the screen. Your configuration changes will be saved, and your Inbox redisplays.

Accessing Gmail Using POP

To check and read your Gmail mail via POP, you must specify the following values when creating an account for receiving incoming mail:

- **Email Address:** *your-username*@gmail.com
- **Server Type:** POP
- **Server:** pop.gmail.com
- **Username:** *your-username*@gmail.com
- **Requires Security/Authentication:** Yes, SSL encryption
- **Port:** 995
- **Password:** *your-gmail-password*

NOTE **You should, of course, replace *your-username* and *your-gmail-password* with the appropriate values for your account when creating a new account in Evolution.**

There are always plenty of other configuration options, such as whether you want messages to be stored on the server and for how long, whether you want to automatically check for messages and how often, whether checking mail at this account should be included in a default mail check, and so on. Different mailers handle this information differently. For example, in Evolution, whether you check mail for an account by default depends on whether an e-mail account is enabled in Evolution's main Accounts dialog. If you are using Kubuntu, this same option is a Kontact configuration item that you set when you are defining an e-mail account for reading e-mail. Either way, you'll probably want to enable this option when first creating and configuring POP access to your Gmail account—you can always disable it after you're sure that everything is working correctly.

If your mail client provides integrated incoming/outgoing accounts (as Evolution does), simply continue with the account definition process. If your mail client stores information about incoming and outgoing mail messages separately (as Kontact does), save these changes and create an account for sending your Gmail mail.

Sending outgoing Gmail mail is done via the SMTP protocol. To send your Gmail mail with your mail client, you must specify the following values when creating an account for sending mail:

- **Server Type:** SMTP
- **Server:** smtp.gmail.com
- **Port:** 587 or 465
- **Requires Security/Authentication:** Yes, TLS or STARTTLS. (If this option isn't available, try SSL encryption.)

Test your new account by using an existing e-mail account to send an e-mail message to your Gmail account, and clicking Send/Receive in Evolution (or File ➪ Check Mail in Kontact, if you're visiting this section from Chapter 10). If you can successfully retrieve mail, you should then check that you can send mail using this account. To do this, create a new message using the icon or the File ➪ New menu, selecting the menu item to send a new mail message.

To see the official Google explanation of how to configure POP access to your Gmail account, see `http://mail.google.com/support/bin/answer.py?answer=13287`. If you have any problems with the explanation provided in this section, you may want to check that site, just in case something has changed since this book was written.

Accessing Gmail Using IMAP

To check and read your Gmail mail via IMAP, you must specify the following values when creating an account for receiving incoming mail:

- **Email Address:** *your-username*@gmail.com
- **Server Type:** IMAP
- **Server:** imap.gmail.com
- **Username:** *your-username*@gmail.com
- **Requires Security/Authentication:** Yes, SSL encryption
- **Port:** 993
- **Password:** *your-gmail-password*

NOTE **You should, of course, replace *your-username* and *your-gmail-password* with the appropriate values for your account when creating a new account in Evolution.**

There are always plenty of other configuration options, such as which folders to check mail in, whether you want to automatically check for messages and how often, whether checking mail at this account should be included in a default mail check, and so on. Different mailers handle this information differently. For example, in Evolution, whether you check mail for an account by default depends on whether an e-mail account is enabled in Evolution's main Accounts dialog. If you are using Kubuntu, this same option is a Kontact configuration item that you set when you are defining an e-mail account for reading e-mail. Either way, you'll probably want to enable this option when first creating and configuring POP access to your Gmail account—you can always disable it after you're sure that everything is working correctly.

If your mail client provides integrated incoming/outgoing accounts (as Evolution does), simply continue with the account definition process. If your mail client stores information about incoming and outgoing mail messages separately (as Kontact does), save your current changes and create an account for sending your Gmail mail.

Sending outgoing Gmail mail is done via the SMTP protocol. To send your Gmail mail with your mail client, you must specify the following values when creating an account for sending mail:

- **Server Type:** SMTP
- **Server:** smtp.gmail.com
- **Port:** 587 or 465
- **Requires Security/Authentication:** Yes, TLS or STARTTLS. (If this option isn't available, try SSL encryption.)

Test your new account by using an existing e-mail account to send an e-mail message to your Gmail account, and clicking Send/Receive in Evolution (or File ➪ Check Mail in Kontact, if you're visiting this section from Chapter 10). If you can successfully retrieve mail, you should then check that you can send mail using this account. To do this, create a new message using the icon or the File ➪ New menu, selecting the menu item to send a new mail message. Make sure that you specify that the identity of the sender is your Gmail account.

To see the official Google explanation of how to configure IMAP access to your Gmail account, see `http://mail.google.com/support/bin/answer.py?answer=78799`. If you have any problems with the explanation provided in this section, you may want to check that site, just in case something has changed since this book was written.

Reading Mail from Other Web-Based Mail Providers

Many free e-mail providers such as Hotmail and its successor, Microsoft Live, Yahoo!, mail.com, and so on provide free e-mail accounts that are designed to be accessed via webmail—in other words, via a browser-based interface. By default, these webmail-only e-mail accounts can therefore not be accessed via standard protocols such as POP or IMAP. Thanks to the cleverness and hard work of zillions of open source software developers everywhere, you can access these accounts via POP if you install and run a local software process that knows how to access these webmail sites, but provides a POP interface that you can talk to. I am a big fan of a package called FreePOPs! (`www.freepops.org`), which does exactly that. Many other open source and commercial packages provide similar functionality to FreePOPs!, but FreePOPs! is available in the Ubuntu repositories, is easy to set up, and has always worked for me. If you're interested in similar packages, you may want to look at YPOPs! (`www.ypopsemail.com/`) or MrPostman (`http://mrpostman.sourceforge.net/`). These are all generic solutions that provide a POP interface for your existing mail reader.

There are a variety of reasons why you may want to access a free webmail account from a mail client such as Evolution or Kontact. Some of my favorites are:

- It's nice to be able to use the same application for all of your e-mail.
- Using a POP client with a webmail account enables you to archive all of your mail on your local machine, which is reassuring for the paranoid among us.
- It's tricky to access webmail when you're offline.
- You also don't have to worry about synchronizing address books using a directed graph that can end up looking like an M. C. Escher drawing if you have enough e-mail accounts.

Note that being able to access a webmail account via POP does not enable you to send mail through that account. As mentioned several times in this chapter, retrieving incoming mail is done via protocols such as POP and IMAP, while outgoing mail is either sent using a remote mail server that you communicate with using SMTP or by contacting a local mail server via a traditional UNIX/Linux Sendmail interface. You have a few choices for sending mail using your webmail account:

- Continue to use the webmail interface to send mail for this account (which is something of a pain).
- Send mail for your webmail account using an existing SMTP server to which you already have access.
- Follow the instructions provided in Bonus Chapter 9 on how to install, configure, and run your own mail server. It's really not that hard, and think how powerful you'll feel!

The next two sections explain how to install FreePOPs! and how to configure your mail client to access a free webmail account through it.

NOTE **Another alternative to using FreePOPs! or another, similar package, is to switch to a different mail reader such as Mozilla's Thunderbird, which provides support for webmail via its Webmail extension (`http://webmail.mozdev.org/`). As a long-time Evolution and KMail, er, I mean, Kontact, fan, that's not an option for me, but your mileage may vary...**

Installing and Using FreePOPs!

The FreePOPs! package is available in Ubuntu's universe repository as the package freepops, and can easily be installed using apt-get, aptitude, Synaptic, or the Adept tools, as explained in Chapter 23. You should probably also install the freepopsdoc package, which will install the FreePOPs! documentation into `/usr/share/doc/freepops-doc`—you can't lose by having all of the documentation for a package.

Once you've installed the freepops package, the FreePOPs! startup file is installed as `/etc/init.d/freepops`. The FreePOPs! server is not started by default. For testing purposes, you can start it manually by executing a command such as `sudo /etc/init.d/freepops restart` (I use `restart` rather than `start` just in case a process is already running for some reason). You should then add this startup script into your system's startup process, typically by creating a symbolic link to the FreePOPs! startup file in the `/etc/rc2.d` directory used to start processes under runlevel 2, which is the default runlevel for Ubuntu and Kubuntu systems. You should also create an equivalent symlink for use when shutting down your system. These commands would look like the following:

```
sudo ln -s /etc/init.d/freepops /etc/rc2.d/S99freepops
sudo ln -s /etc/init.d/freepops /etc/rc2.d/K99freepops
```

For more detailed information about startup and shutdown scripts and the Ubuntu/Kubuntu startup and shutdown process, see Chapter 22, "Understanding the Ubuntu Startup and Shutdown Processes."

After you've run the FreePOPs! startup script, the FreePOPs! daemon will be running on your system, which you should be able to see in a process listing. You are now ready to configure an e-mail account so that you can read your webmail via POP, as described in the next section. Even though I've been using this package for quite a while, I still find it incredibly cool that I'm able to do this. Thanks, open source community!

Creating POP Accounts for Use with FreePOPs!

After installing FreePOPs!, you'll need to create accounts that enable you to access it. To check and read your webmail via POP, you must specify the following values when creating an account for receiving incoming mail:

- **Email Address:** *your-username*@*host*.com
- **Server Type:** POP
- **Server:** localhost
- **Username:** *your-username*@host.com
- **Requires Security/Authentication:** No
- **Port:** 2000
- **Password:** *your-password*

In these examples, replace "host" with "hotmail," "live," "mail," "yahoo," or whatever other domain name is used by your free webmail host. Similarly, you should replace *your-username* and *your-password* with the appropriate values for your account when creating a new account in Evolution.

If your mail client provides integrated incoming/outgoing accounts (as Evolution does), simply continue with the account definition process, specifying one of the SMTP servers discussed in the introductory text for the section entitled "Reading Mail from Other Web-Based Mail Providers." If your mail client stores information about incoming and outgoing mail messages separately (as Kontact does), save the current changes for POP access to your webmail account and create an account for sending the mail associated with your webmail account.

Test your new account by using an existing e-mail account to send an e-mail message to your webmail account, and clicking Send/Receive in Evolution (or File ➪ Check Mail in Kontact, if you're visiting this section from Chapter 10). If you can successfully retrieve mail, you should then check that you can send mail using the outgoing mail server that you associated with this account. To do this, create a new message using the icon or the File ➪ New menu, selecting the menu item to send a new mail message. Make sure that you specify that the identity of the sender is your webmail account.

NOTE **POP/POP3 access to Yahoo! mail accounts is supposedly available to Yahoo! mail account holders in a non-USA Yahoo! domain. (I haven't verified this.) In the United States, you have to upgrade (and therefore pay money) for the Yahoo! Mail Plus package in order to get a send/receive POP e-mail account at Yahoo!. See `http://mailplus.mail.yahoo.com` for more information.**

TIP **To ask questions about or make suggestions for FreePOPs, see the FreePOPs user forums at `http://freepops.diludovico.it`.**

Managing Contacts

Contacts are the people with whom you often interact and exchange e-mail, and about whom you therefore want to maintain contact information: their full name; their surface address for snail mail; their e-mail address(es) at home, work, or both; their home page; any Instant Messaging addresses they have; their corporate or academic affiliation; their home, office, and mobile telephone numbers; and so on. Collections of contact information are typically referred to as an address book.

Maintaining complete and easily accessed contact information is an important part of any PIM application. Evolution provides integrated support for managing local contact information, as well as the ability to contact networked address books via standard protocols. The next sections explain how to configure and use Contacts information within Evolution.

Configuring Evolution's Contact Support

Evolution provides two basic levels of contact customization: configuring how Evolution's mail composer interacts with the information in your contact database(s), and configuring the source of your contact information. Unless you're using Evolution in a commercial or academic environment, you probably won't care about remote sources of contact information, so I'll focus on configuring contact/mail interaction first, and explain using contact information from remote LDAP servers second, for those of us who are lucky enough to be able to do so.

Configuring Contact/E-mail Interaction

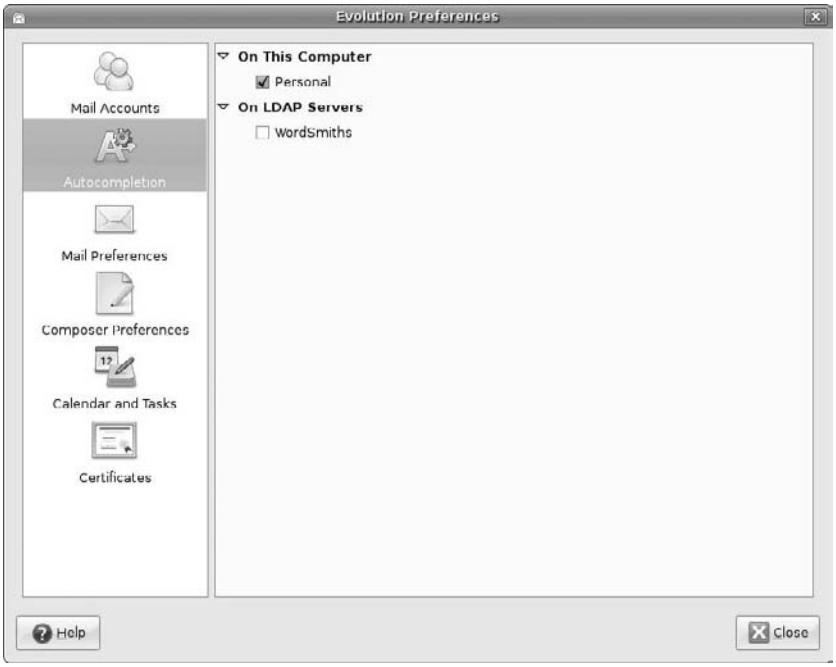
Evolution's Contacts doesn't require much independent configuration, but how an Evolution e-mail account interacts with a Contacts database, and which one it interacts with, can be configured in different areas of the standard Evolution Preferences dialog. Select Edit ➪ Preferences (or press Shift+Ctrl+S) in Evolution to display this dialog.

By default, Evolution will automatically attempt to look up names and e-mail addresses from your default address book as you type them when composing mail, which is known as autocompletion because Evolution is automatically trying to provide complete names or e-mail addresses based on the characters that you have typed so far. When the characters that you have typed match one or more names or e-mail addresses in any of your address books, a list of all possible completions will be displayed in the appropriate field (To, CC, BCC) in Evolution's composition window. You can accept any of these potential completions by using the mouse to select that name or e-mail address, or continue typing to further refine the list of possible matches.

The default address book for Evolution is the Personal address book located on the computer that you are using to compose mail. If you have multiple address books, such as an LDAP address book that you can access via a network, you will have to activate these address books as other sources that Evolution should check when trying to automatically complete names and e-mail addresses. To do this, click the Autocompletion icon in the left pane of the Evolution Preferences dialog, which displays a dialog like the one shown in Figure 9-28.

FIGURE 9-28

Activating Autocompletion in available address books



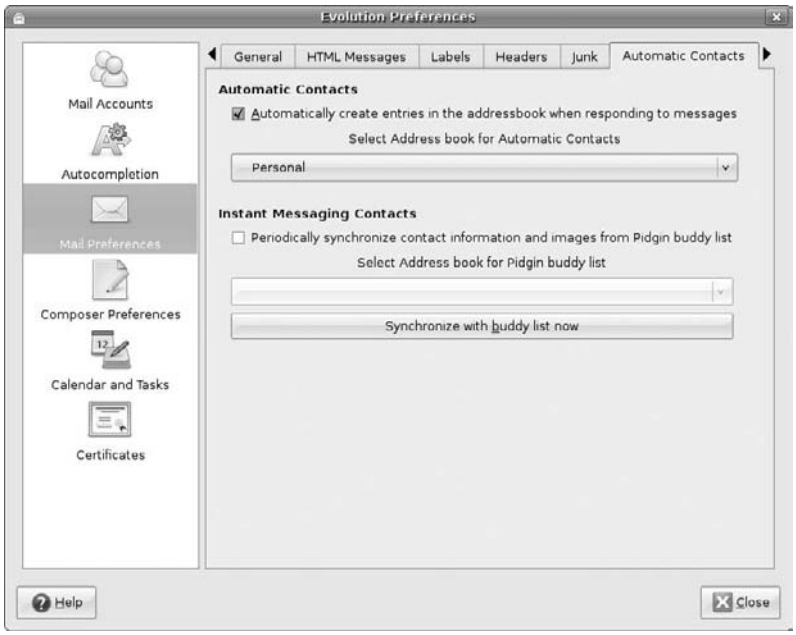
Address books that will be used for autocompletion are checked. To add another address book, simply select the checkbox to the left of its name. Once you're done making changes, click Close to close the Preferences dialog and return to Evolution, or click another icon to configure a different aspect of Evolution.

One of the features that I find most useful when building an address book is Evolution's ability to automatically create contacts based on the e-mail addresses in any messages that you respond to. Whenever I receive mail from an old (or new) friend, that friend is added as a contact to my Personal address book automatically when I reply to the message. No more looking for a message from someone in order to figure out his or her e-mail address. Evolution simply does the right thing—once this configuration option is enabled, of course.

Activating automatic contact creation is done on the Mail Preferences pane of the Evolution Preferences application. To access related settings, first open the Preferences dialog by selecting Edit ➪ Preferences (or pressing Ctrl+Shift+S) in Evolution, and clicking the Mail Preferences icon in the left pane of the Evolution Preferences dialog. Next, select the Automatic Contacts tab at the far right of the top of this dialog. This displays a screen like the one shown in Figure 9-29.

FIGURE 9-29

Configuring Automatic Contact creation and synchronization



To automatically create contacts for any recipient of e-mail that you send, click the checkbox to the left of the aptly named "Automatically create entries in the addressbook when sending message" option. Once this checkbox is selected, you must also select the address book to which you want to add those entries. To add them to the address book on the computer that you are using to compose and send mail, click the "Select Address book for automatic contacts" drop-down list and select Personal. If you have more than one local address book, or have access to a networked LDAP address book, you can select any of these, but you can identify only a single address book to which contacts will be added automatically.

If you use Pidgin, the instant messaging client discussed in Chapter 13, the dialog shown in Figure 9-29 also enables you to automatically pull contact information from your Pidgin buddy list. To activate this, click the checkbox to the left of the Periodically synchronize…entry and click the drop-down list to select the Pidgin address book that you want to pull from. After selecting an address book, you can click the Synchronize with buddy list now button to extract current contact information from your Pidgin buddy list and copy it into Evolution. Don't worry—this information will still be available in Pidgin. This is strictly a non-destructive pull operation.

Accessing Remote Contact Information via LDAP

When using Ubuntu and Evolution on a home computer system, Evolution's local address book provides all of the power that you'll ever need to use to create and manage information about friends, family, and other contacts. However, Evolution's address book can also be seamlessly integrated with business or academic settings that maintain central, networked sources of contact information. If you work for a small company, you may know (or be able to figure out) everyone's e-mail address. When you work for IBM, that's a bit trickier, hence the need for centralized address books. Most networked address book information is provided by LDAP (Lightweight Directory Access protocol) servers, which are well-supported in Evolution.

To add an LDAP server to the sources of address book information that are consulted when you compose mail and identify addressees, click New in Evolution's icon bar and select Address Book from the drop-down list that appears. You can also create a new address book by right-clicking on any existing Address Book and using the menus by selecting the File ➪ New ➪ Address Book menu item. This displays the dialog shown in Figure 9-30.

FIGURE 9-30

Creating a new address book in Evolution



Once this dialog displays, click the Type drop-down list and select On LDAP Servers. A dialog like the one shown in Figure 9-31 displays.

In the Name field, enter the name that you want to associate with the LDAP server that you want to contact, and which you will see in the LDAP Servers list. (This does not need to be the hostname of the LDAP server.) Next, enter the hostname or IP address of the LDAP server in the Server field, and change the port if necessary (port 389 is the default). Next, specify the type of security required to connect to and retrieve information from the LDAP server. This is either No encryption, TLS encryption (Transport Layer Security), or SSL encryption (Secure Sockets Layer) connection. Finally, select the Login method required to authenticate to the server, which is one of Anonymously, Using email address, or Using distinguished

name (DN). If you select one of the latter two, you must enter the appropriate information—either your e-mail address, or your distinguished name in your LDAP domain, which is typically something like `CN=vonhagen,OU=Users,DC=ldap,DC=vonhagen,DC=org` (which means "the Users\vonhagen entry in the LDAP server running on the host ldap.vonhagen.org").

FIGURE 9-31

Creating an entry for an LDAP server



If searches for LDAP entries begin at a specific LDAP directory (such as users in the previous example), you can specify this by clicking the Details tab and entering the LDAP starting point in the Search base field.

Finally, click OK to save your LDAP address book definition. You should now be able to look up and retrieve e-mail addresses for people from the LDAP address book.

TIP **If Evolution does not display an error message but does not seem to retrieve remote contact info from your LDAP server, you may need to restart your evolution-data-server process. This is the Evolution process that communicates with remote services such as LDAP. To identify this process, execute a command such as the following:**

```
$ ps ax | grep evolution
3014  ? Sl 5:34  evolution --component=mail
3021  ? Sl 0:02 /usr/lib/evolution/evolution-data-server...
20031 ? Sl 0:00 /usr/lib/evolution/2.12/evolution-exchange...
```

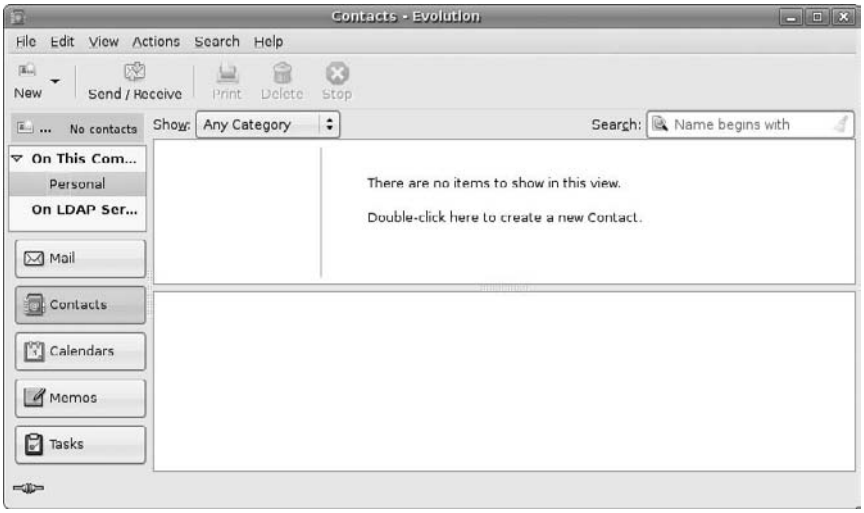
In this case, you will need to restart process 3021, which you should be able to do with the command `kill -HUP 3021`. Worst case, terminate all of your evolution processes using the `kill -9 PID` command (where PID is the process ID of each process), and restart evolution, which will restart all necessary processes.

Adding and Editing Contact Information

To view and edit your contacts, click the Contacts button in the lower-left corner of the main evolution window. A screen like the one shown in Figure 9-32 appears.

FIGURE 9-32

Your Initial Contacts view



Unless you imported contact information into Evolution during the Evolution setup process or have already sent mail in Evolution with the "Automatically add to addressbook" option active (described earlier in this chapter in the section "Configuring Contact/E-mail Interaction"), the first time that you view your contacts, you won't see any because you have not yet defined any. You can type Ctrl+N in the Contacts view to create a new contacts entry. This displays a dialog like the one shown in Figure 9-33.

This dialog enables you to enter all of the information that you can (or want to) about the contact that you are defining. The Contact tab, displayed by default when you first add a new contact or edit an existing one, contains name, e-mail address, phone, and instant messaging information for a contact. This dialog also provides two other tabs:

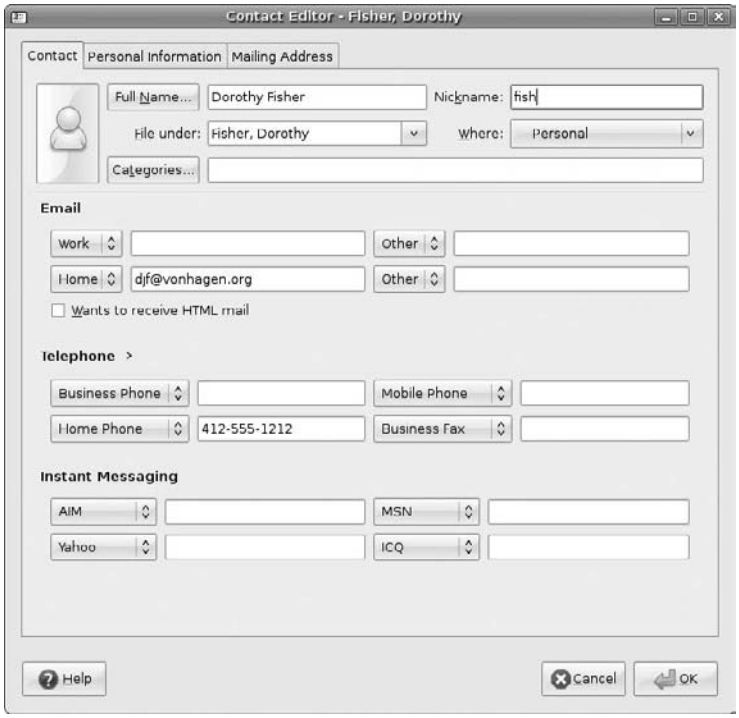
- **Personal Information:** Enables you to record information about a contact's Web presence (home page, online calendar, blog, and so on), professional information such as their position, employer, manager, assistant, and so on), and personal information such as the name of their significant other, their birthday, anniversary, and so on.
- **Mailing Address:** Enables you to enter their personal and work e-mail addresses, and a spare address field for things such as the location of their beach house or hunting camp.

Once you've entered whatever information you want to record about this contact, click OK to close the contact information dialog and return to the general contacts view. Once you've defined one or more

contacts, the Contacts view will look something like Figure 9-34, with summary information about the currently selected contact displayed in the bottom half of the Contacts pane.

FIGURE 9-33

Defining contact information



To see this summary information about any defined contact, select that contact from the Contacts list pane (above the summary information). You can scroll left to right through this list to locate any existing contact. Once you've accumulated a large number of contacts, you may find it more convenient to search for a contact by entering a portion of their name in the Search area above the contacts list and pressing Return to view any matches for that string. You can also click the icon at the left of the search box to search by e-mail address or by the contents of any field in the contact information.

To edit or view detailed information about any existing contact, double-click on the entry for that contact to display again the dialog shown in Figure 9-33, which enables you to correct or add information about that contact. If you were just checking information, click Cancel to close this dialog. If you have made changes, click OK to save the updated information.

After viewing or editing information, you can return to Evolution's mail view at any time by clicking the Mail button in the list in the lower-left corner.

FIGURE 9-34

Summary information about a selected contact



Working with vCards

VCards are text files with a `.vcf` file extension that contain contact information. The format of these files was designed to provide a standard method of exchanging contact information between users regardless of the e-mail client or computing platform that they use.

If you receive e-mail containing a vCard, you can easily import it into Evolution's contacts list by scrolling down in Evolution and clicking Save in Address Book.

If Evolution doesn't give you this option for some reason, you can also save vCards to text files on your system by right-clicking the attachment and selecting Save As. You can then use Evolution's File ➪ Import command to import the vCard into your address book.

If you want to share your Evolution contact information with someone else as a vCard, you can create a vCard for a single contact or save all of your contacts to a single vCard file. To create a vCard file containing your selection(s), select one or more contacts and right-click on the selected items. Select Save as vCard. A dialog displays that enables you to specify the name of the file to which you want to save your selection(s). To simply mail the vCard version of your selection(s), right-click on the selection(s) and select Forward Contact. A new mail message opens in a composition window with a vCard containing your selection(s) as an attachment.

Defining Mailing Lists

Many of us belong to Internet mailing lists, Google groups, Yahoo! groups, and other focused e-mail discussions of our favorite topics, including perhaps the various Ubuntu mailing lists discussed in Chapter 1. These types of mailing lists are usually handled by specialized software packages such as MailMan (`www.gnu.org/software/mailman/`), which handles things such as keeping the list members anonymous, providing a single e-mail address to send and respond to mail, and supporting online archives of mailing list traffic. While these requirements are almost mandatory for Internet-wide mailing lists, they are overkill for smaller mailing lists, such as your family, friends with a similar sense of humor, former colleagues from a previous employer, and so on. For most of us, having a simple way of associating multiple e-mail addresses with a single list name would suffice—which is exactly what Evolution provides.

Evolution makes it easy to set up your own mailing lists, known as Contact Lists, for personal communications by defining a Contact List name and then associating specific people from your contacts address book with that list. You can then send mail to the name of the list, and Evolution handles sending that identical message to all of the people on that list.

You can create a new Contact List in various ways:

- Pressing Ctrl+Shift+L anywhere in Evolution
- Selecting the File ➪ New ➪ Contacts List menu entry
- Clicking the New button and selecting Contacts List from the drop-down menu
- Right-clicking in the Contact List when in the Contacts view, and selecting New Contact List from the context-sensitive menu

Performing any of these actions displays the dialog shown in Figure 9-35.

FIGURE 9-35

Creating a new Contact List

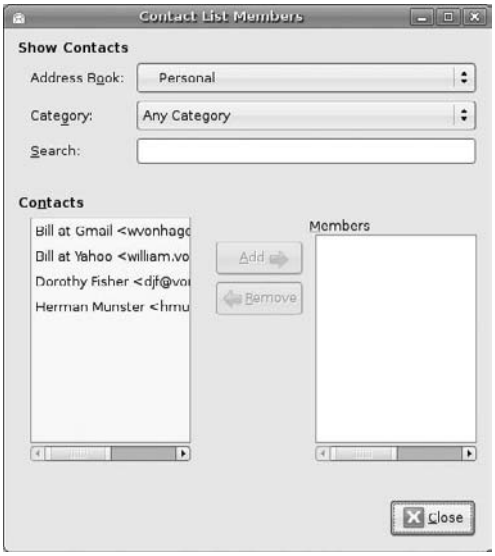


In the List name field, enter the name of the list that you want to create. You can then add e-mail addresses to the new list in any of the following ways:

- Click Select to display the dialog shown in Figure 9-36, in which you can select existing Contact entries and click Add to add them to the current Contact List. When you're done, you can click OK to return to the dialog shown in Figure 9-35.

FIGURE 9-36

Adding e-mail addresses to a Contact List



- Drag and drop contact entries into the Members area from the pane in which you contacts are listed.
- Manually enter e-mail addresses that you want to add to the list, pressing Return after each address.

Once you've created and populated a Contact List, you can configure whether or not the e-mail that you send to the list contains all of the e-mail addresses of its members. The "Hide addresses when sending mail to this list" checkbox is selected by default. When this checkbox is selected for a mailing list, the recipients of your messages will see Undisclosed Recipients as the name to which the message was sent, and will not be able to reply to anyone other than to you (because the names of the other list members are hidden). If you deselect this checkbox, all of the e-mail addresses of the list members will be shown in the To: field of your mail message, and any recipient can therefore reply to every recipient by selecting the Reply To All command.

NOTE **Hiding the recipients of mail to a Contact List is done by putting their e-mail addresses in the BCC (Blind Carbon Copy) field. Evolution displays a warning message when sending a message that has no visible recipients, to which you can simply respond OK if you only wanted to send mail to list members without letting them know who else received the message.**

To save your new Contact List and close the dialog shown in Figure 9-35, click OK. You can now send mail to the list by entering the name of the list in a new mail message, just as you would with a standard e-mail address. You can also send mail to the list by right-clicking on its name in your list of contacts in the Contacts view and selecting Send Message to List from the context-sensitive menu that displays.

Managing Your Calendar

After e-mail and managing contact information, the last critical component of a personal information management application is the ability to manage one's calendar. Evolution's Calendar component, accessed by clicking the Calendars button in the lower-left corner of the Evolution GUI, is both easy-to-use and powerful, making it quite easy for you to track and manage your personal calendar on your personal machine, creating custom reminders for the items on your calendar. By default, a local calendar named Personal is created in every Evolution installation.

Evolution's calendar support isn't limited to maintaining your calendar on your machine—you can publish information about your calendar to a Web site for anyone to see, and you can also connect to remote calendaring resources via standard protocols such as WebCal and CalDAV to know what others are doing or for group scheduling purposes. Because of this dual focus on individual and group calendaring items, Evolution's calendar module handles defining personal appointments and group meetings slightly differently.

The next few sections discuss the highlights of configuring your Evolution calendar, and explain how to create local calendar entries with reminders, configure your Evolution calendar to integrate remote calendaring resources, and share your free/busy information with others.

Basic Calendar Configuration

You can configure a variety of aspects of the Evolution Calendar. To configure the Evolution Calendar, select the Edit menu's Preferences command and selecting the Calendars and Tasks icon from the left pane of the Preferences dialog. This displays the dialog shown in Figure 9-37.

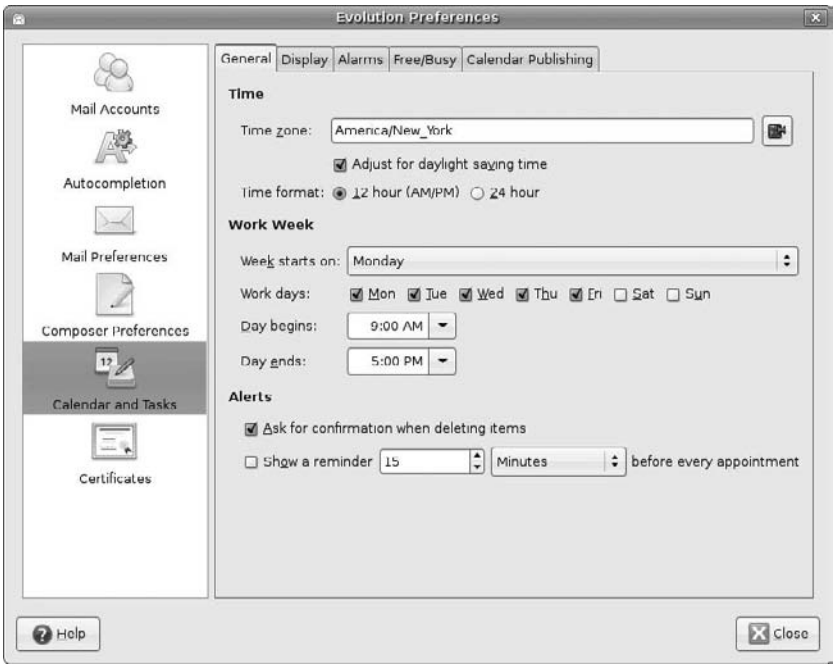
Evolution enables you to configure the time zone, time format (12 hour or 24 hour), the days of the week and hours of the day of your work week, the minimum size of a schedulable appointment, and many options having to do with how appointments, meetings, and alarms are displayed. Evolution's Calendar module provides a good set of default values, and many people do not need to customize the configuration of their calendar. However, I've tweaked a few things over the years, so I'll highlight those in the following subsections.

Defining Your Work Day and Work Hours

By default, Evolution assumes a work week of Monday through Friday, and working hours of 9 a.m. to 5 p.m. Not all of us are lucky enough to work those hours, so you may want to change this to reflect your actual schedule. You can do this from the screen shown in Figure 9-37, by selecting the days of the week that you work, and selecting the hours on which your work day begins and ends from the appropriate drop-down lists.

FIGURE 9-37

General Calendar configuration items in Evolution



Configuring your work week and hours doesn't affect Evolution's ability to schedule items at any time, but configuring the hours in your work day is necessary before you can print your calendar for a given day. By default, any single-day calendar that you print will contain only entries that fall within your workday.

After changing any values on this tab, click Close to close the Evolution Preferences dialog. The new work week and work hours that you selected will automatically be displayed on your calendar.

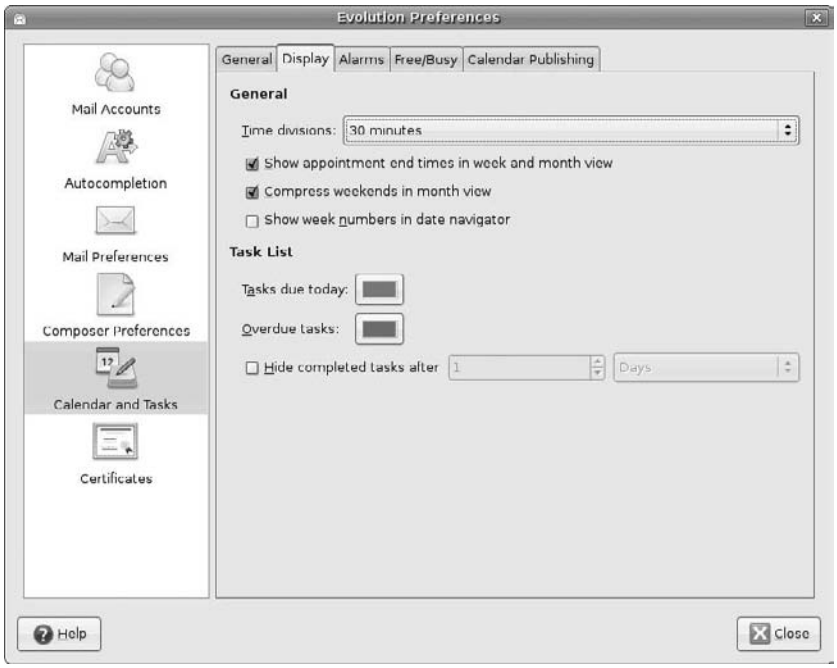
Setting Event Granularity on Your Calendar

Another commonly modified Evolution calendar setting is the time divisions that are shown on your calendar. Evolution's default calendar setting uses 30-minute time divisions, which means that you can't schedule an appointment for a time period less than 30 minutes. If, like me, you prefer to schedule things in 15-minute increments, you can change the time divisions by selecting the Display tab in the Calendar and Tasks configuration dialog, which is shown in Figure 9-38.

Once this dialog displays, you can select a different time division, such as 15 minutes, from the Time divisions drop-down list. After changing this value, click Close to close the Evolution Preferences dialog. The new time divisions that you selected will automatically be displayed on your calendar.

FIGURE 9-38

Configuring Calendar granularity and display in Evolution



Creating and Editing Personal Calendar Entries

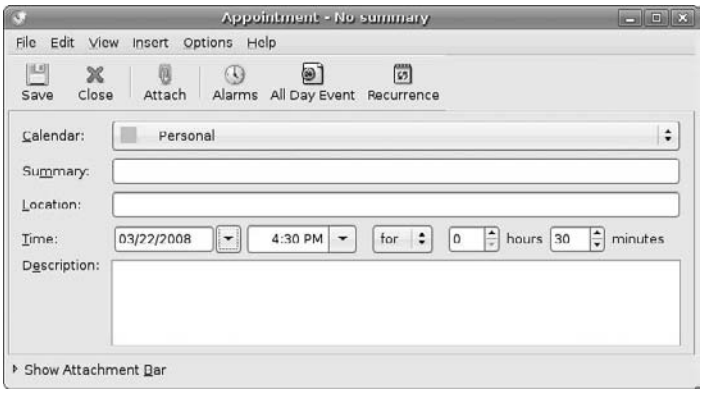
Evolution's calendar module enables you to define two different types of events in your personal calendar: appointments that are scheduled for a specific day and time (and which can repeat), and All Day Appointments that are scheduled for a specific day. Like most things in Evolution, you can create these types of entries for your personal calendar in a number of different ways:

- Pressing Ctrl+Shift+A anywhere in Evolution to create a personal appointment
- Selecting the File ➪ New ➪ All Day Appointment or Appointment menu entries
- Clicking the New button and selecting All Day Appointment, or Appointment from the drop-down menu
- Right-clicking in a calendar when in the Calendar view, and selecting All Day Event or Appointment from the context-sensitive menu

Performing any of these actions displays a dialog like the one shown in Figure 9-39. This dialog is for defining an appointment, not an all-day appointment, so it includes information about the start and stop time for the appointment. The all-day appointment dialog is almost identical, except that it does not include a start and stop time.

FIGURE 9-39

Defining a new appointment



Enter the basic information about the appointment that you want to schedule by entering a Summary, which is the basic description of the appointment that will display on your Evolution calendar. (Keep it brief but intuitive.) You can optionally enter a location for the appointment, and enter a more verbose description of the appointment (including meeting notes, phone numbers, and so on) in the Description field. You should also double-check the start and stop times for the appointment and correct them if necessary. (If you created your appointment by right-clicking on your calendar and selecting New Appointment, the default start time for the meeting will reflect the time division in which the cursor was located when you selected that command.)

After filling in the details of your appointment, you may want to set an Alarm by clicking the Alarms icon in the Appointment dialog. This will cause a pop-up reminder dialog to display on your screen at the time for which the alarm was set.

Clicking the Alarms icon displays a dialog like the one shown in Figure 9-40.

FIGURE 9-40

Defining an alarm for an appointment



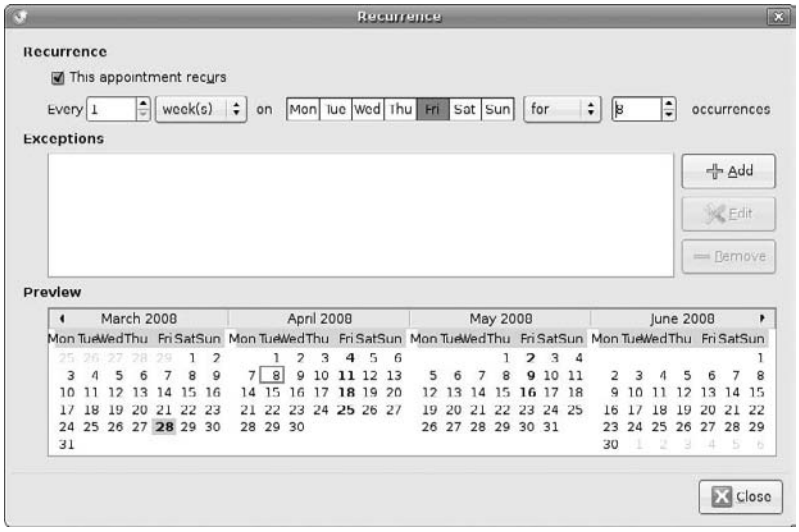
Click the Alarm field to show the drop-down menu of possible alarm times shown in Figure 9-40. After selecting an alarm, it is displayed in the alarm list window. You can then click Close to close this dialog and return to the appointment dialog, or select Customize from the Alarm list to customize the text for the alarm that you selected.

When an alarm goes off, a dialog displays on your screen, reminding you of the upcoming appointment. You can either click Close to terminate the alarm, or click Snooze to dismiss the alarm for the specified period (the default is 5 minutes), after which the alarm dialog will re-display. if you need a bit more time, you can also customize the snooze period on the alarm dialog before clicking Snooze.

After defining an appointment and setting any alarms that you want to use as reminders, you can also configure this appointment to be a recurring appointment. This means that the appointment will be automatically added to your calendar based on the period that you specify for its recurrence. To configure an appointment as a recurring appointment, click the Recurrence icon in the toolbar for the appointment that you are defining. This displays a dialog like the one shown in Figure 9-41.

FIGURE 9-41

Defining recurring appointments



Once this dialog displays, click the This appointment recurs checkbox and use the drop-down lists to specify when it re-occurs and for how long. If there are specific days on which you will not be having this appointment, such as on holidays, you can select those days on the calendar previews and click Add to add them to the list of exceptions.

Once you have finished defining any recurrence of the appointment that you are creating, click Close to return to the Appointment dialog. You can then click Save to save the current state of this appointment and continue working on it, or click Close to close this dialog and return to Evolution's Calendar view.

After you have defined an appointment, you can edit it at any time by double-clicking on the appointment in Evolution's calendar view.

Creating and Editing Meetings

Meetings differ from appointments and all-day appointments in that they involve multiple people. Entries for the meetings that you define will be sent to all of the participants that you define for the meeting, and will show up on their calendars if they accept the invitation.

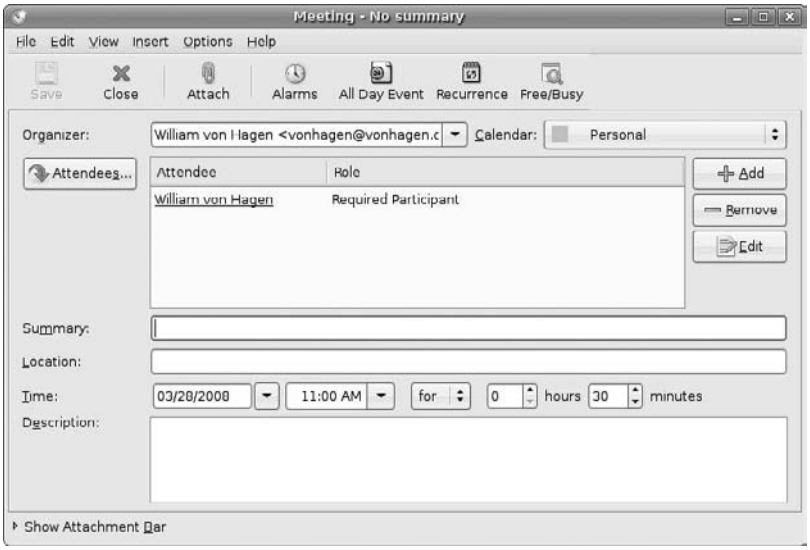
You can create a new meeting in a number of different ways:

- Pressing Ctrl+Shift+E anywhere in Evolution
- Selecting the File ➪ New ➪ Meeting menu entry
- Clicking the New button and selecting Meeting from the drop-down menu
- Right-clicking in a calendar when in the Calendar view, and selecting Meeting from the context-sensitive menu

Performing any of these actions displays the dialog shown in Figure 9-42.

FIGURE 9-42

Creating a new meeting

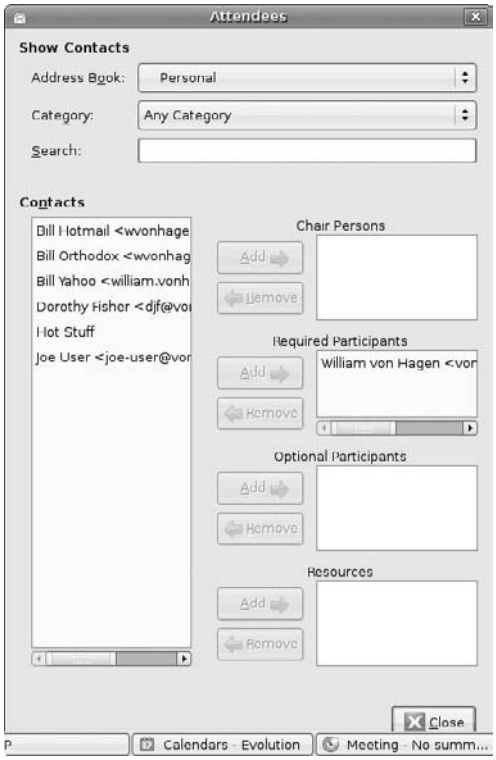


Providing basic information about your meeting, setting alarms, and defining recurring meetings are all done as described in the previous section. The primary difference between creating an appointment and

a meeting is that you will probably want to invite others to the meeting. You can add a single attendee by clicking Add and entering the name or e-mail address of that person. However, it is much easier to add other people to your meeting by clicking the Attendees button to display the dialog shown in Figure 9-43.

FIGURE 9-43

Adding attendees to a meeting



This dialog enables you to select a contact and then click Add beside the different dialogs that identify their role in the meeting: Chair Persons, Required Participants, and Optional Participants. If your conference rooms are listed in an LDAP address book, you can also reserve them by adding them as resources.

Once you have finished identifying the people and resources associated with this meeting, click Close to return to the Meeting dialog. You can then click Save to save the current state of this meeting and continue working on it, or click Close to dismiss this dialog and return to Evolution's Calendar view.

After you have defined a meeting, you can edit it at any time by double-clicking on the meeting in Evolution's calendar view.

Integrating Online Calendars

Being able to view other people's calendars is an important capability in today's networked business and academic environments. Tracking your own time is convenient, but being able to see other people's calendars can save you a significant amount of time when trying to meet with someone or simply track them down. To support viewing other people's calendar's, the Evolution Calendar module supports the most common types of online calendars and associated protocols that are in use today: Google calendars, CalDAV, and WebCal. Evolution also supports weather calendars, which enable you to display current temperature and general weather condition information in Evolution's calendar view.

Integrating an online calendar is done by creating a new calendar of the appropriate type in Evolution and filling in any information that is required to connect to that calendar. You can create a new calendar in Evolution in a number of different ways:

- Selecting the File ➪ New ➪ Calendar menu entry
- Clicking the New button and selecting Calendar from the drop-down menu
- Right-clicking on an existing calendar in the Calendar Summary window at the upper left of the Calendar view, and selecting New Calendar from the context-sensitive menu

Performing any of these actions displays the dialog shown in Figure 9-44.

FIGURE 9-44

Creating a new calendar



Click the Type drop-down list and select the type of calendar that you want to define. Your choices are Google, CalDAV, On This Computer (a new local calendar that is located only on your machine), On The Web (a WebCal calendar), and Weather.

Next, specify a name for this calendar. This name can be anything that is meaningful to you, but most people tend to use a name that identifies the source of the calendar information.

Finally, depending on the type of calendar that you are creating, you will have to enter information so that Evolution's Calendar module knows how to contact the calendar. The information required for each type of calendar is as follows:

- **Google:** Enter your Google username. This provides read-only access to your Google calendar.
- **CalDAV:** Enter the URL of the calendar that you are contacting and the username to use when contacting the calendar, if necessary
- **On This Computer:** No additional information is required, because this is simply a local calendar.
- **On The Web:** Enter the URL of the WebCal or iCal calendar that you want to access, whether or not a secure connection is required, and the username to use when contacting the calendar, if necessary.
- **Weather:** Click the location button to display a dialog that enables you to browse through available locations and select the one for which you want to display weather information. You can click Units to specify whether you want weather information in Celsius or Fahrenheit, and use the Refresh selector to specify how often you want to update the weather information.

Once you have entered the appropriate information for the type of calendar that you are creating, click OK to close the New Calendar dialog and save the configuration of your new calendar.

After having created a calendar, you can edit its settings at any time by right-clicking on its name in the Calendar Summary window at the upper left of the Calendar view, and selecting Properties from the context-sensitive menu.

TIP **A number of predefined WebCal calendars are available on the web at sites such as `http://icalshare.com`. Two of my favorites are the Lord of the Rings calendar available at `webcal://boobah.com/calendars/lotr.ics`, and the Pittsburgh Pirates season calendar at `webcal://sports.yahoo.com/mlb/teams/pit/ical.ics`. (Other teams are available from similar Yahoo! URLs, but I don't care about any of them.)**

Publishing Your Free/Busy Information

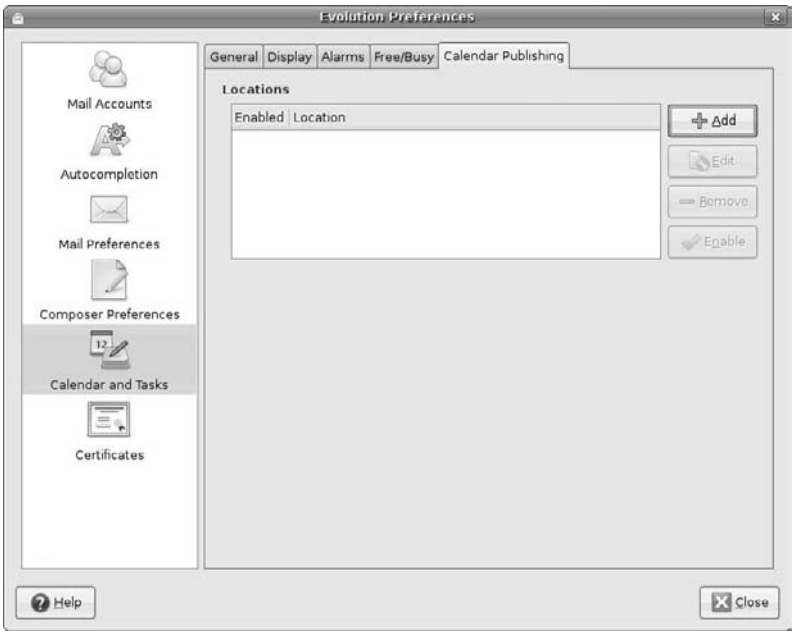
The previous section discussed how to connect to remote calendars over a network. Google's online calendar is available from anywhere, and protocols such as iCal are fine if a business or academic site that you're involved with has a shared calendaring system that you can access from anywhere. However, even if you are only managing your own personal calendar, it can be quite useful to export your calendar and/or free/busy information so that others can access it directly.

Evolution makes it easy to publish your personal calendar in iCal or Free/Busy formats. I do this for my wife when I'm traveling on business, so that she has some idea of what it is that I'm doing all day and when we may be able to actually talk rather than simply exchanging dueling e-mail. Publishing this information yourself simply writes this information to a file that others can then read to see your calendar information. I publish my calendar in iCal format to a file on my Web site, which my wife then loads into Evolution in order to see my updated information.

Publishing your calendar is configured using the Evolution Preferences dialog, which you can display by selecting Edit ➪ Preferences. Once this dialog displays, click Calendar and Tasks, and then click the Calendar Publishing tab. This displays a dialog like the one shown in Figure 9-45.

FIGURE 9-45

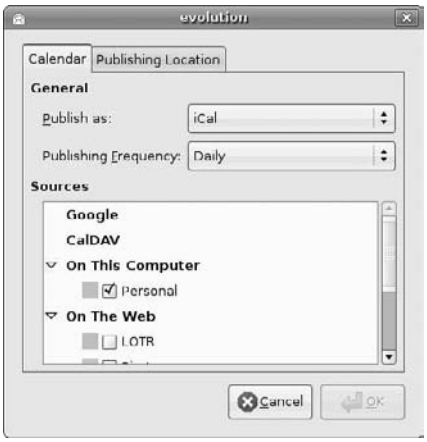
Configuring Calendar Publishing in Evolution



To define a publishing schedule and format for your calendar, click Add. This displays a dialog like the one shown in Figure 9-46.

FIGURE 9-46

Specifying publishing format and frequency for a calendar



Select the checkbox beside the name(s) of the calendar(s) that you want to publish, and select the format in which you want to publish your calendar by clicking the Publish As drop-down list and selecting either iCal or Free/Busy. The former provides a detailed listing of your calendar, while the latter simply shows when you are busy. Finally, select the frequency with which you want to publish your calendar information. Choices are Daily, Weekly, and Manually, by selecting the Actions menu's Publish Calendar Information command. Once you've made these selections, click the Publishing Location tab. The tab shown in Figure 9-47 displays.

FIGURE 9-47

Specifying publishing connection and location information



Select the protocol that you want to use from the Service Type drop-down list. Supported protocols are SSH, Public FTP (no authentication required), FTP (with login), Windows Share (using the CIFS protocol), WebDAV and Secure WebDAV, and Custom (which enables you to specify a different protocol).

Next, enter the hostname or IP address of the server to which you want to publish your calendar, and the name of the file that you want to save it to. You can then enter an (optional) custom port number and any username and password that you need to use to connect to the remote site. Being able to enter a custom port is quite convenient because many sites move well-known services to different ports in order to try to hide them from hackers who expect to find specific services listening on specific, traditional ports.

After making these changes, click OK to save your calendar publishing information. You can then test whether you've configured everything correctly by selecting the Actions menu's Publish Calendar Information command, and then checking the remote location to make sure that a file containing your calendar information was correctly created there.

Additional Sources of Information About Evolution

Don't get your hopes up, Charles Darwin fans! This section discusses other sources of information about the Evolution PIM/e-mail application, not the theory of evolution.

Because the Evolution application is generally very popular and is the default e-mail application on most GNOME-based systems (such as Ubuntu), there are a good number of Web sites that can provide you with additional information about using Evolution, as well as mailing lists that you can join to ask questions, file problem reports, and so on. Some of my favorites are as follows:

- **Evolution Forums (`http://nabble.com/Gnome---Evolution-f1297.html`):** If you prefer to use a forum interface instead of a mailing list, Nabble provides a great collection of forums that present the various Evolution mailing lists in a user-friendly forum format. Messages posted to these forums go directly to the associated mailing list.
- **Evolution IRC Channel:** If you're a fan of Internet Relay Chat and have a client handy (such as the XChat-GNOME IRC client that is provided with Ubuntu), you can discuss developing and using Evolution in the #evolution channel on irc.gimp.org. You should lurk for a little while before broadcasting your questions, just to get a feel for the channel and to reduce your chances of being flamed if you ask "the wrong question" or one for which an answer is readily available in the documentation or elsewhere.
- **Evolution Mailing Lists:** You can subscribe to a general mailing list for Evolution users (known as the *evolution-list*—catchy, no?) at `http://mail.gnome.org/mailman/listinfo/evolution-list`, and you can also subscribe to a more advanced Evolution hackers list (evolution-hackers) at `http://mail.gnome.org/mailman/listinfo/evolution-hackers`. For general information about all GNOME-related mailing lists, see `http://mail.gnome.org/mailman/listinfo/`.
- **Evolution Project Home Page (`http://gnome.org/projects/evolution`):** This is the home page of the Evolution project and, as such, is the definitive starting point for looking for additional information about Evolution, the source code for the latest version of Evolution, and so on.
- **Evolution Wiki (`http://go-evolution.org/Main_Page`):** This is a great site for information about the current release of Evolution and for insights into (or making suggestions for) the capabilities that will be present in upcoming versions of Evolution.

Any reasonable Internet search engine will display many more sites with information about Evolution, but these are the primary official conduits for Evolution information.

Summary

This chapter explained how to configure and use Evolution, a powerful e-mail client and Personal Information Manager, focusing on the e-mail aspects of Evolution. Evolution is a great e-mail client that supports all popular e-mail protocols and provides unrivaled connectivity to specific types of mail servers such as Microsoft Exchange e-mail servers. Reading and sending e-mail is one of the most common tasks that anyone does with their personal computer nowadays. Equally common is surfing the Web, which is explained in the next chapter. To satisfy your Web surfing requirements, Ubuntu automatically installs the Firefox Web browser, the most powerful, extensible, and just plain fun Web browser available today.

Chapter 10

Managing E-mail and Personal Information with Kontact

After Web browsers, the program you use to send and receive e-mail is the Internet's killer application. As with Web browsers, there are many different client applications for sending and receiving e-mail, especially in the open source community. GNOME-based Ubuntu systems typically use Evolution, the e-mail application and PIM (Personal Information Manager) suite that was discussed in Chapter 9. KDE-based systems such as Kubuntu provide a personal information management suite called Kontact that is made up of a number of existing and familiar KDE applications.

Kontact provides all of the capabilities that people normally expect from a modern PIM application by combining KMail (the traditional KDE e-mail client), Akregator (an RSS feed aggregator), KAddressBook (a contact management utility), KNode (a Usenet news reader), KNotes (an application for note taking), KOrganizer (a scheduling/calendar application), and KPilot and KitchenSync (Palm handheld synchronization applications) into the single Kontact "meta-application." All of these applications mesh seamlessly thanks to the magic of the KDE KPart framework, which provides both a consistent application development model and a mechanism for developing applications as components that can subsequently be integrated into interoperable application suites. Although all of the applications that make up Kontact were originally standalone applications (and can still be used as such), why not take advantage of the power of an integrated suite of related applications whenever possible?

NOTE **Throughout the rest of this section, I'll refer to the application that we're running as Kontact, regardless of which subcomponent we're actually exploring.**

IN THIS CHAPTER

- Configuring e-mail accounts
- Sending and receiving e-mail
- Organizing and filtering e-mail
- Configuring spam and virus detection
- Creating and managing contacts
- Using the calendar

Starting Kontact

You can start Kontact in a variety of ways:

- By selecting the Kontact icon from the Quick Launcher portion of the kicker

- By selecting Kontact – Personal Information Manager from the K Menu's Internet menu
- By executing the `kontact` command from any command-line environment, such as a Konsole or xterm window

Starting Kontact from the Quick Launcher is the fastest and easiest way to start the program. When you first start Kontact, a window like the one shown in Figure 10-1 displays.

FIGURE 10-1

Kontact's Summary View window



This window is known as Kontact's Summary view because it provides a summary of pending calendar and to-do items for the current date, lists any special dates that it knows about which occur within the next week, and provides a summary of pending e-mail in certain folders. Above this summary information, Kontact displays the current system date and the name associated with the current login, which it obtains from the Linux password file. A vertical row of icons along the left edge of the Summary View window provides single-click access to the different applications that are a part of Kontact.

The Summary View window is interesting, but it won't provide any meaningful information unless you configure Kontact for at least one e-mail account and enter contacts, calendar items, to-do items, and so on.

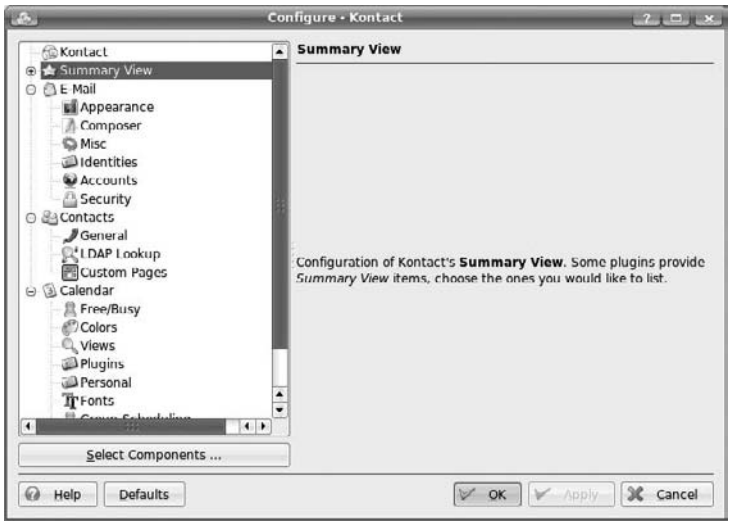
Configuring Kontact

To configure Kontact, select the Settings ➪ Configure Kontact menu item. This displays the configuration dialog shown in Figure 10-2.

If you're just getting started with Kontact, the first thing that you'll probably want to do is to set up an e-mail account, as explained in the next few sections. Reading and sending e-mail in Kontact and configuring various aspects of how e-mail is displayed and organized in Kontact are discussed later in this chapter.

FIGURE 10-2

Kontact's configuration dialog



Configuring E-mail in Kontact

Kontact can receive mail from mail servers that speak POP3 and IMAP and can also be set to work with local and maildir format mailboxes that are populated by fetchmail or similar Linux applications. For outgoing mail, Kontact can communicate with any Simple Mail Transport Protocol (SMTP) server or any server that is compatible with the standard Linux/UNIX sendmail application.

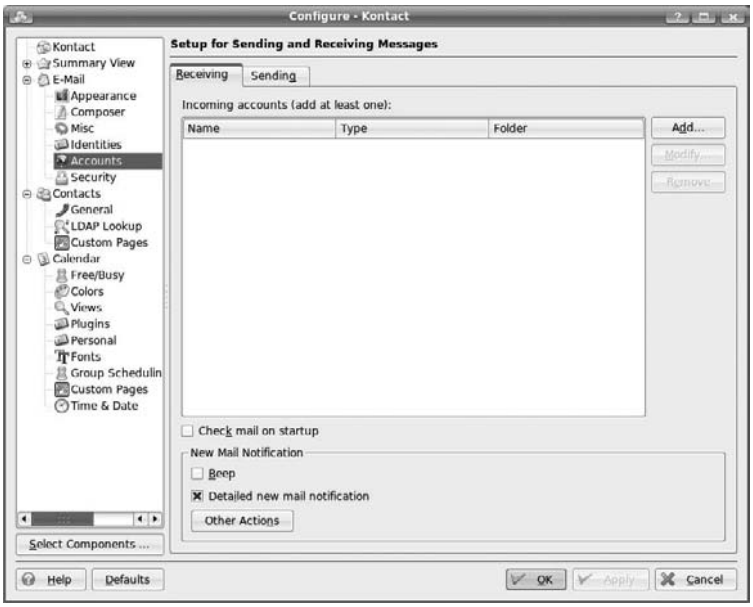
When getting started with Kontact, your first step is to create an account that you can use to create, send, and receive e-mail. An account is simply the collection of settings associated with a specific e-mail address. The first account you will want to define is the one associated with your ISP or a Web-based account (Gmail, Yahoo!, Hotmail, and so on) through which you currently send and receive your personal e-mail. If you have more than one e-mail address—say, a Gmail address and a Yahoo! address—you can create separate accounts for each of these later. For now, let's simply configure your personal e-mail account.

NOTE **You can configure the e-mail component of Kontact (KMail) in two ways—either through Kontact, by selecting the Settings ➪ Configure Kontact menu item when you are in Kontact's Summary view, or when running KMail, by selecting the Settings ➪ Configure KMail menu item. Both of these display configuration dialogs that are identical, with the exception that Kontact's configuration dialog displays all of the configuration settings for all Kontact applications in a pane at the left of the dialog. All of the configuration examples in this chapter use Kontact's configuration dialog in order to stress that Kontact unifies a number of different KDE applications into a single PIM suite.**

To set up an e-mail account in Kontact, select the E-mail ➪ Accounts menu item from the tree in the left pane of the configuration dialog shown in Figure 10-2. A screen like the one shown in Figure 10-3 displays, listing any e-mail accounts that you may already have defined and enabling you to define new ones. (As you can see, the example in Figure 10-3 shows that no e-mail accounts have been set up yet.)

FIGURE 10-3

The initial account creation screen in Kontact



NOTE **If you send and receive e-mail through a Web-based account, such as Gmail or Hotmail, you will need to configure the account that you use to retrieve mail based on the characteristics of those sites. Gmail supports POP3 and IMAP, but services such as Microsoft's Hotmail or Live sites and Yahoo! support Web-based e-mail only. Fortunately, some amazing open source tools enable you to access Web-based mail just as though the sites used real e-mail protocols. See the section of Chapter 9 entitled "Accessing Free and Web-Based E-mail" for detailed information about how to access Gmail via POP3 or IMAP, and how to access free or Web-based e-mail using any mail client that supports POP3 (as Kontact does).**

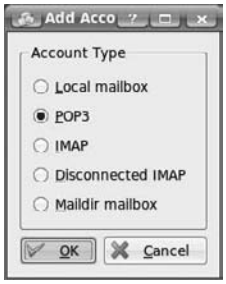
To create a new account, click Add. A dialog like the one shown in Figure 10-4 displays, allowing you to specify the protocol used to retrieve e-mail or the location of a specific mailbox that Kontact should use to locate and store e-mail.

You will probably want to select either POP3 or IMAP. Most mail servers nowadays speak POP3, which is an enhanced version of the Post Office Protocol (POP) that adds some extensions to standard POP mail server capabilities and is the version of POP supported by Kontact. (POP3 is a superset of POP, so POP3 "includes" POP support.) The second most popular protocol nowadays is IMAP, which also includes Microsoft Exchange mailers because they speak a variant of IMAP. The primary difference between POP3 and IMAP is where your mail is permanently stored. POP/POP3 mailers download entire mail messages to the e-mail client and attempt to minimize long-term mail storage on the remote server. IMAP mailers permanently store messages on the remote server and initially download only message headers to the e-mail client, retrieving the body of an e-mail message (and any attachments) only when the user explicitly requests it. Most ISPs prefer POP/POP3 mailer servers because using them reduces disk space requirements at their end, because messages are typically downloaded to your computer and then deleted

from the mail server once they are read (although you can change this behavior, as explained later). POP/POP3 mail accounts are usually local to a specific computer system for this reason. IMAP mail servers are generally preferred by businesses or other entities whose users need to be able to access e-mail from multiple computer systems and locations. Even when POP/POP3 mail is stored in a network filesystem that is shared between multiple systems, such filesystems are not always accessible from remote locations. (For additional comparative information about POP3/IMAP mailers, see the sidebar "A Quick Guide to POP vs. IMAP" in Chapter 9.)

FIGURE 10-4

Specifying an e-mail retrieval protocol or mailbox location



NOTE **The Local mailbox and Maildir mailbox options are for users who are using only local mail service on a single machine, or who use an application such as fetchmail to automatically retrieve e-mail from a mail server and deposit it into a local mail storage location. These are somewhat advanced topics, so I won't discuss them here. If you think that you might need to use one of these options, contact your system administrator.**

Select the protocol you use to communicate with the mail server for this account and click OK. (I have my own domain and associated POP3 mail server, so I selected POP3.) If you are configuring a POP3 account, you'll next see a dialog like the one shown in Figure 10-5.

If you are configuring an IMAP account, you'll see a dialog like the one shown in Figure 10-6.

When first opened, these dialogs display the General tab for the Modify Account dialog. The mandatory fields on this tab are as follows:

- **Account name:** This is an arbitrary name that you want to use to identify this e-mail account. People often use the name of the location through which this account sends and receives e-mail (your domain name, Gmail, and so on) or the e-mail address associated with this e-mail account (for example, vonhagen@vonhagen.org or wvonhagen@gmail.com). The name you use is completely up to you.
- **Login:** This is the login name of the individual that is retrieving e-mail from the mail server. Typically this is just a login name (vonhagen, wvonhagen, and so on), but some ISPs that use a centralized mail server to handle the e-mail needs of many different users and domain require a login of the form username%domain-name. In my case, if necessary, this would be something like vonhagen%vonhagen.org.

- **Password:** This, of course, is the password you've previously established in order to send and receive e-mail.
- **Host:** This is the fully qualified name of the mail server that hosts the e-mail for this account. This is typically mail.domain.name (for example, mail.vonhagen.org).
- **Port:** This is the port through which e-mail is retrieved for the mail retrieval protocol associated with this account. For POP/POP3 mailers, this is port 110 by default. For IMAP mailers, this is port 143 by default. This value should already have been filled in based on the retrieval protocol you specified using the dialog shown in Figure 10-4.
- **Identity:** This is the name of the Kontact identity associated with this e-mail account. Multiple identities are useful if you use e-mail from different sources, enabling you to send e-mail using the identities associated with these different accounts. The default value, Default, specifies the identity with which you send mail if you don't specify a particular identity. This value is all you need if you have only one e-mail account. (You will have to set the e-mail address associated with that identity before you can actually send mail, as explained later in this section.)

FIGURE 10-5

Configuring a POP3 e-mail account in Kontact

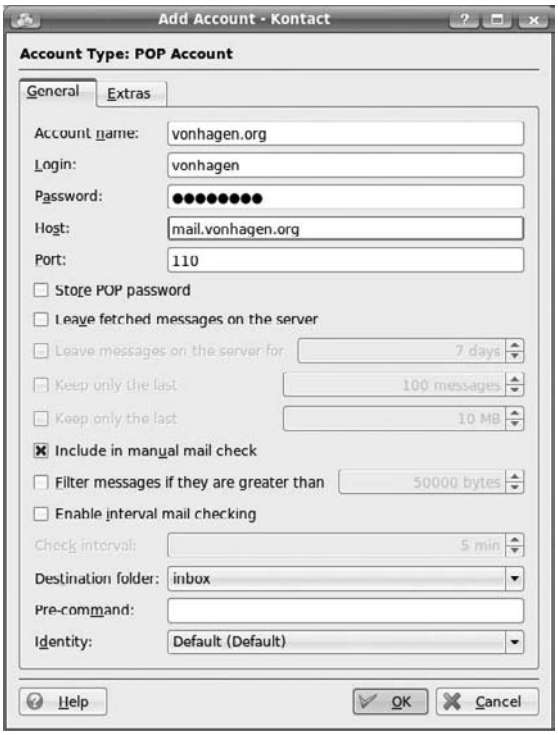
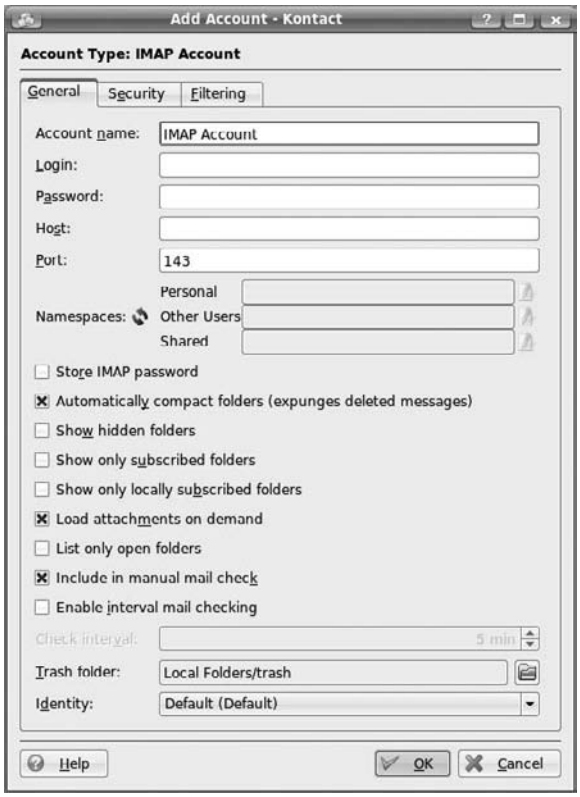


FIGURE 10-6

Configuring an IMAP e-mail account in Kontact



In addition to these mandatory items, you can, if you wish, set a number of options that differ based on whether you are defining a POP3 or IMAP e-mail account. The next two sections discuss protocol-specific options for POP3 and IMAP accounts.

Configuring POP3 Accounts for Receiving Mail

The previous section explained the configuration details that apply to any e-mail account you create in Kontact, regardless of whether you are receiving mail via the POP3 or IMAP protocol. This section explains the additional items shown in Figure 10-5 that are specific to POP3 accounts.

For POP3 accounts, you can specify the following:

- **Store POP password:** This option enables Kontact to remember the POP password for this account so that you don't have to supply it each time you check mail. Your password is stored in the KWallet application, which is a secure password management system for the KDE environment. KWallet is analogous to the Mac OS X Keychain or the GNOME Keyring.
- **Leave fetched messages on the server:** This option causes Kontact to leave messages on the server as it reads them. This option is useful if you want to be able to check messages from a

number of machines, but actually store them on only a single central machine. For example, I use this option when traveling and checking mail from my laptop, so that once I get home, I can still download the messages to the central system on which I "officially" read mail, on which I do backups, and so on.

- **Leave messages on the server for (period):** This option is enabled only if you've selected the Leave fetched messages on the server option; it enables you to leave messages on the server for a set period of time. Messages still on the server whose age exceeds this value are deleted the next time that you read mail. This option is useful to occasionally clean up old messages that have been read, but which may not have been correctly deleted from the server.
- **Keep only the last:** These two options (with the same name) are enabled only if you've selected the Leave fetched messages on the server option; they enable you to delete messages on the server if the number of saved messages exceeds a specific number of messages, or if the size of your mailbox on the server exceeds a set limit. I've never used the first of these options, keeping only a specified number of messages, but the second is convenient when checking mail on systems that have a set limit on the amount of storage that you can use. Setting the "by size" value to something under this limit helps you guarantee that your e-mail storage requirements won't exceed that limit.
- **Include in manual mail check:** This option determines whether this account will be included in those that are checked when you use Kontact's Check Mail or Synchronize commands.
- **Filter messages if they are greater than:** This option enables you to automatically exclude messages that exceed a specified size (50K is the default). This can protect you from automatically receiving e-mail messages with huge attachments that might exceed your storage requirements. I rarely activate this option because I often expect (and get) e-mail from various publishers with large attachments.
- **Enable interval mail checking:** This option enables you to specify a time period at which the account will automatically be checked for new messages. This can be handy, but it can also be a distraction if you get lots of mail but are actually trying to get work done on your machine.
- **Destination folder:** This option enables you to set the name of the folder in which incoming e-mail will be stored. A default value of Inbox is provided, but you can change this to the name of any other mail folder that exists in Kontact. Creating and using mail folders is discussed later in this chapter in the section "Organizing and Filtering Mail."
- **Pre-command:** This option enables you to specify the full pathname of a command that should be run before Kontact checks for mail. For example, you can run a script to retrieve remote e-mail and forward it to an SMTP server or stash it in a local location (such as a maildir or mbox file). The current account can then examine the local location or query a specific mail server.
- **Identity:** This option enables you to specify the identity (typically, a username and an e-mail account) that is associated with this account when mail is sent. Kontact provides a default identity to send mail from but also enables you to define as many other identities as you want. There is typically a one-to-one mapping between the number of identities and the number of accounts you've defined. You must specify an e-mail address for at least your default identity before you can send mail.

After setting the options you're interested in and that apply to your POP3 account, you must still do the following:

- Configure any security mechanisms used when retrieving e-mail
- Configure the account you're using for sending mail through your POP3 account

- Configure any security mechanisms used when sending e-mail
- Define at least one identity for use when sending mail

The next section discusses details for configuring IMAP accounts. Given that you are reading this, and thus probably have just configured a POP3 account, you may want to skip to the section "Defining Security for Retrieving E-mail." The other topics are explained in subsequent sections.

TIP **Even if you receive mail from multiple POP3 mail hosts, as I do, you will still have only a single set of local mail folders into which to receive mail. In these cases, I create a separate folder to hold incoming mail from each POP3 account, giving each of them a name that is evocative of the POP3 mail host. I then set the Destination folder for each account to the new folder, and add that folder to Kontact's Favorite Folders pane. Creating and using mail folders is discussed later in this chapter in the section "Organizing and Filtering Mail."**

Configuring IMAP Accounts for Receiving Mail

The beginning of this section explained the configuration details that apply to any e-mail account that you create in Kontact, regardless of whether you are receiving mail via the POP3 or IMAP protocols. This section explains the additional items shown in Figure 10-6 that are specific to IMAP accounts.

For IMAP accounts, you can specify the following:

- **Store IMAP password:** This option enables Kontact to remember the IMAP password for this account so that you don't have to specify it each time you check mail for this account in Kontact. Your password is stored in the KWallet application, which is a secure password management system for the KDE environment. KWallet is analogous to the Mac OS X Keychain or the GNOME Keyring.
- **Automatically compact folders (expunges deleted messages):** Enabling this option causes Kontact to automatically delete messages on the server when you leave a folder in which you have deleted messages.
- **Show hidden folders:** Enabling this option causes Kontact to display folders to which you are not subscribed. This can be handy if you are also accessing public folders on an IMAP server, and are not sure what folders may be available in addition to those to which you are currently subscribed.
- **Show only subscribed folders:** Enabling this option causes Kontact to display only folders to which you are currently subscribed on the server. If this option is selected, Kontact will query the server for the list of subscribed folders each time you check mail for this account.
- **Show only locally subscribed folders:** Enabling this option causes Kontact to display only folders to which you are locally listed as being subscribed.
- **Load attachments on demand:** Enabling this option causes Kontact to retrieve information about attachments when checking e-mail, but it will only download an attachment from the mail server when you explicitly open it in Kontact (usually by clicking on the attachment).
- **List only open folders:** Choosing this option checks for new mail and updated only in folders that are currently open in Kontact. This can save time when updating IMAP mail listings over slow connections or on slow systems.
- **Include in manual mail check:** This option determines whether this account will be included in those that are checked when you use Kontact's Check Mail or Synchronize commands.

- **Enable interval mail checking:** This option enables you to specify a time period at which the account will automatically be checked for new messages. This can be handy, but can also be a distraction if you get lots of mail but are actually trying to get work done on your machine.
- **Trash folder:** This option enables you to set the name of the folder in which deleted mail will be stored. A default value of Trash in your local IMAP folders is provided, but you can change this to the name of any other mail folder that exists in Kontact. Creating and using mail folders is discussed later in this chapter in the section "Organizing and Filtering Mail."
- **Identity:** This option enables you to specify the identity (typically, a user name and an e-mail account) that is associated with this account when mail is sent. Kontact provides a default identity for the person who is sending mail, but also enables you to define as many other identities as you want. There is typically a one-to-one mapping between the number of identities and the number of accounts that you have defined. You must specify an e-mail address for at least your default identity before you can send mail.

After setting the options that you are interested in and which apply to your IMAP account, you must still do the following:

- Configure any security mechanisms used when retrieving e-mail
- Configure the account that you are using for sending mail through your IMAP account
- Configure any security mechanisms used when sending e-mail
- Define at least one identity for use when sending mail

These topics are explained in the next few sections.

Defining Security for Retrieving E-mail

Specifying the password that you use to authenticate to a POP3 or IMAP server is only one aspect of accessing your e-mail in a secure fashion. Both POP3 and IMAP can support a wide range of authentication mechanisms, so you must also configure Kontact to know the authentication method that is used by the mail server that you are contacting. Similarly, many mail servers support different mechanisms for retrieving your e-mail in a secure fashion, which prevents people from simply capturing the network traffic (also known as "packets") that flow between your e-mail client and the remote mail server in order to read your mail.

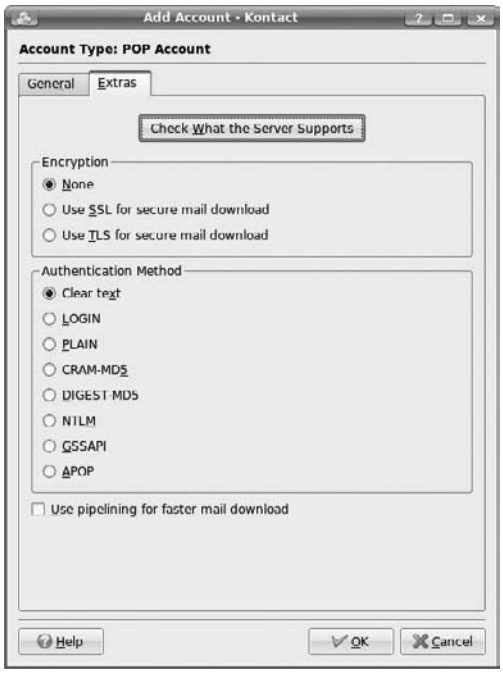
To define the authentication and encryption mechanisms used by the mail server from which you are retrieving e-mail, click the Extras tab shown in Figure 10-5 or the Security tab shown in Figure 10-6. Selecting either of these tabs displays the dialog shown in Figure 10-7, on which you can set these options.

Unless you're a mind reader or your ISP has better documentation than most, figuring out the encryption and authentication methods supported by a remote mail server can be truly painful. Luckily, Kontact provides a button labeled Check What the Server Supports, which you can click to query the mail server that you identified on the General tab, and which will select the correct values for you on this screen.

Once these values are displayed, click OK to close the Modify Accounts dialog, and proceed to the next section to continue configuring Kontact.

FIGURE 10-7

Defining authentication options for retrieving e-mail



Defining Accounts for Sending E-mail

After you have created one or more accounts for receiving e-mail in Kontact, you must also configure the flip side of those accounts, which are the accounts used for sending mail through the same ISP or other Internet site. As discussed earlier in this section, POP/POP3 and IMAP are the most popular protocols used to retrieve e-mail. The equivalent configuration items for sending mail are SMTP (Simple Mail Transport Protocol) mail servers located elsewhere on a network and Sendmail mail servers that are running on the machine that you are currently using. Sendmail is not a protocol, but is the name of a specific type of mail server that uses an extended SMTP command syntax. As discussed in Bonus Chapter 9, many other Linux mail servers are available in addition to Sendmail, but Sendmail is so popular (and common) on UNIX and Linux systems that most of these provide their own program called sendmail that behaves like the actual Sendmail mail server.

To create an account for sending e-mail, select the Sending tab (see Figure 10-3) and click Add. A dialog like the one shown in Figure 10-8 displays.

Select the appropriate transport protocol for the mail server that you are using for sending mail, and click OK. Selecting Sendmail on this dialog therefore means that you have a mail server running on the machine where Kontact is running, and that Kontact can communicate locally with this mail server. In this case, you will see the Add Transport dialog shown in Figure 10-9, in which you can specify the name and full pathname of the local sendmail-compatible application.

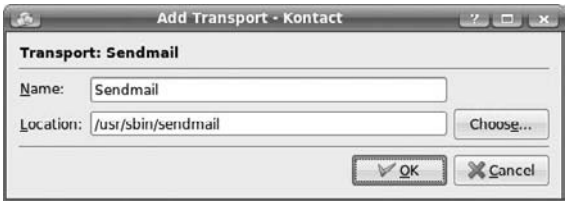
FIGURE 10-8

Selecting the transport protocol used for sending mail



FIGURE 10-9

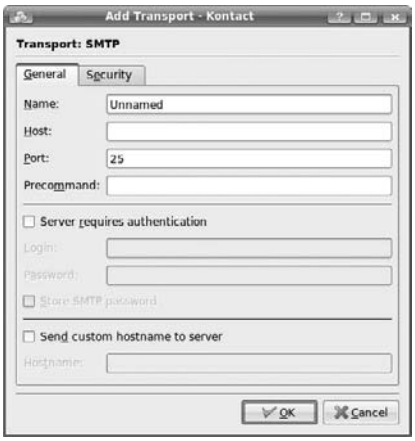
Configuring a local Sendmail server used to send e-mail



If you select SMTP (which is far more common), the version of the Add Transport dialog shown in Figure 10-10 displays.

FIGURE 10-10

Configuring the SMTP server used to send e-mail



The Add Transport dialog's General tab displays when you are providing information about an SMTP mail server. This dialog enables you to specify the following options for the SMTP mail server through which you will be sending e-mail for a given mail account:

- **Name:** This is an arbitrary name that you want to use to identify this outbound e-mail account. People often use the name of the location through which this account sends/receives e-mail (your domain name, Gmail, and so on) or the e-mail address associated with this e-mail account (for example, vonhagen@vonhagen.org or wvonhagen@gmail.com). The name that you use is completely up to you.
- **Host:** the Fully qualified name of the SMTP mail server through which e-mail for this account is sent. This is typically smtp.domain.name (for example, smtp.vonhagen.org).
- **Port:** The port through which e-mail is sent to this mail server. This is traditionally port 25, but many sites use different port numbers to cut down on hacking attempts.
- **Precommand:** Enables you to specify the full pathname of a command that should be run before sending mail.
- **Server requires authentication:** A checkbox that identifies whether the mail server requires authentication in order to send mail through it. Most do, nowadays. If you activate this checkbox, you must specify your login name and password on the mail server. You can also select the Store SMTP password checkbox to cause Kontact to save this password in KWallet.
- **Send custom hostname:** Enables you to specify the hostname (the name of the local machine) that will be sent to the remote server. This provides a workaround for machines that only identify themselves with a short name by default ("foo," as opposed to "foo.vonhagen.org").

Once you have specified these settings, proceed to the next section to provide authentication and encryption information.

Defining Security for Sending Email

Just like authenticating to a mail server when retrieving e-mail, specifying the password that you use to authenticate to an SMTP or local Sendmail server is only one aspect of sending e-mail in a secure fashion. When sending mail via a local Sendmail server, you don't need to specifically authenticate to the mail server because you are already logged in on the machine, and are thus authenticated. If you are configuring Kontact to talk to a local Sendmail server, you can just skip to the next section, "Defining Your Identity."

If you are configuring how you will be sending mail through a networked SMTP mail server, you must also configure Kontact to know the authentication method that is used by the mail server that you are contacting. Similarly, you must also specify the mechanism used for sending your e-mail in a secure fashion, in order to prevent people from capturing the network traffic between your e-mail client and the remote mail server so that they can read the e-mail that you're sending to your mom.

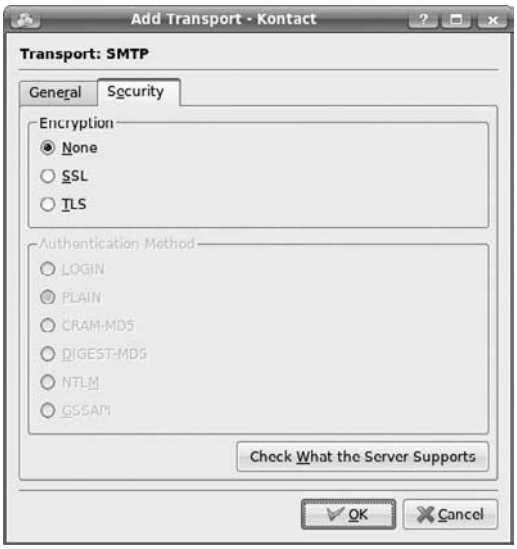
To define the authentication and encryption mechanisms used by an SMTP mail server through which you are sending e-mail, click the Security tab shown in Figure 10-10. This displays the dialog shown in Figure 10-11, in which you can set these options.

Unless you're a mind reader or your ISP has better documentation than most, figuring out the encryption and authentication methods supported by a remote mail server can be truly painful. Luckily, Kontact provides a button labeled Check What the Server Supports that you can click to query the mail server that you identified on the General tab, and which will select the correct values for you on this screen.

Once these values are displayed, click OK to close the Add Transport dialog, and proceed to the next section to continue configuring Kontact.

FIGURE 10-11

Configuring an SMTP server used to send e-mail



Defining Your Identity

As mentioned earlier, you cannot send mail from Kontact until you have defined at least one identity to whom that mail will be sent. When you first start Kontact, it creates a default identity from the full name of the user that is currently logged in. However, Kontact has no way of knowing the e-mail address that is associated with that user because no e-mail accounts have been defined when Kontact is first started. You must edit this default identity and supply an associated e-mail account before you can send mail using Kontact. If you have multiple e-mail accounts, you can subsequently create other identities that you can use when sending mail to ensure that the From and Reply-To e-mail addresses in outgoing e-mail match the e-mail account that you wanted to use.

To define an identity in Kontact, select the E-mail ➪ Identities menu item from the left pane of the configuration dialog shown back in Figure 10-2. The Manage Identities dialog displays, as shown in Figure 10-12, which lists the default identity that was created by Kontact, plus any other identities that you may have already defined.

To finish configuring your default identity or edit any other existing identity, select that identity in the Manage Identities dialog and click Modify. The Edit Identities dialog displays with its General tab selected, as shown in Figure 10-13. This dialog enables you to define an identity or modify an existing one.

You must fill in the Your Name and E-mail address fields. The Organization field is optional, but can be useful if you want to associate yourself with a specific company, academic institution, domain, and so on.

FIGURE 10-12

The initial identity definition screen in Kontact

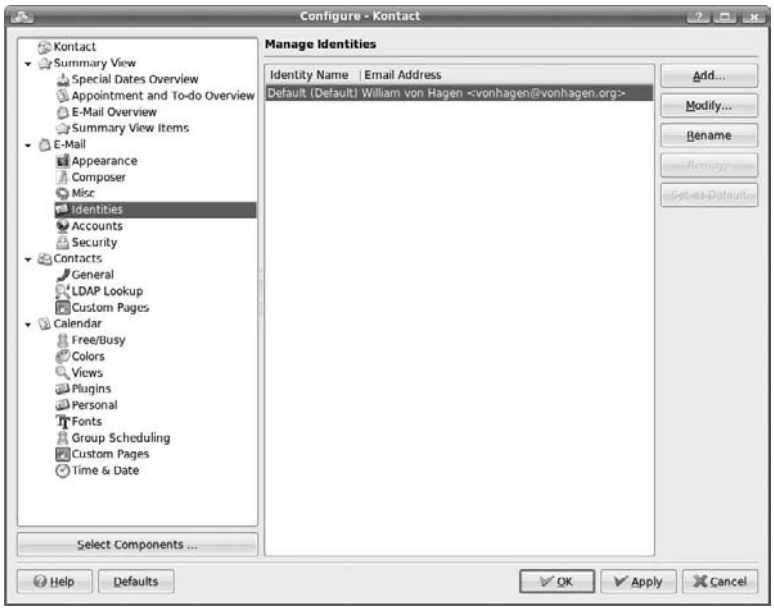
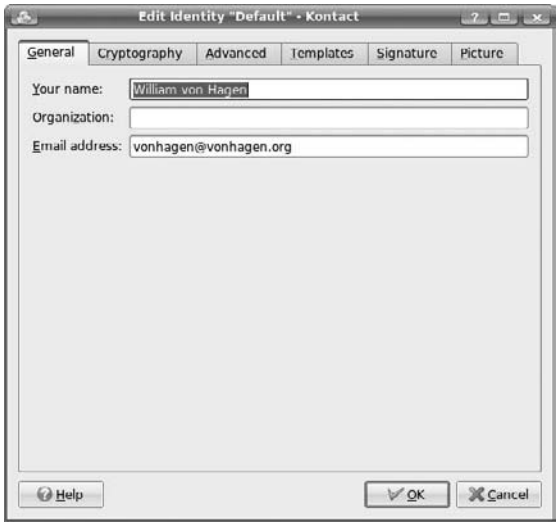


FIGURE 10-13

Defining or updating an identity



The other tabs on the Edit Identity dialog are also optional, but enable you to define and customize the following identity-related items:

- **Cryptography:** Enables you to identify OpenPGP (Open Pretty Good Privacy) and S/MIME (Secure Multimedia Internet Mail Extension) signing and encryption keys/certificates, and to specify the cryptographic message format that you want to use by default for this identity.
- **Advanced:** Enables you to customize aspects of a mail that is created and sent using this identity, such as the Reply-To address (by default, this is the same as the identity's e-mail address); an address to which BCC (Blind Carbon Copy) mail should always be silently sent; the dictionary used to spell-check messages; and the folders in which copies of mail that has been sent, drafts of in-progress messages, and mail message templates for this identity should be stored, and so on.
- **Templates:** Enables you to create a custom e-mail template that will be used by default for all mail sent as this identity.
- **Signature:** Enables you to specify a signature that should be appended to all e-mail sent as this identity.
- **Picture:** Enables you to specify a picture that should always be appended to all e-mail sent as this identity. (This essentially requires that you are using HTML mail.)

Once you have filled in the mandatory values on the General tab and updated any optional values on any other tab, click OK to save your changes and return to the Manage Identities dialog.

To create a new identity, click Add on the screen shown in Figure 10-12. A dialog appears in which you can provide a name for the identity that you are creating. This can be different than the user name that you will supply on the Edit Identity dialog's General tab, shown in Figure 10-13. When creating an identity, the first name that you specify is the name of the identity that will appear in the left column of the Manage Identities dialog shown in Figure 10-12.

After entering a name for the new identity, you can select the appropriate radio button to create the new identity with no predefined settings (With empty fields), use settings information from the Control Center to pre-populate certain fields (Use Control Center settings), or base the default values for the new identity on another existing identity (Duplicate existing identity). Once you have made any radio button selection that you want, click OK in the screen shown in Figure 10-12 to begin configuring the new identity using the dialog shown in Figure 10-13. At this point, the process is the same as that of editing an existing identity, as described earlier in this section.

You've now completed the Kontact settings that are required for sending and receiving e-mail. You'll probably want to test these first to ensure that everything is actually working correctly, but that shouldn't take long. Congratulations in advance!

Configuring Contacts

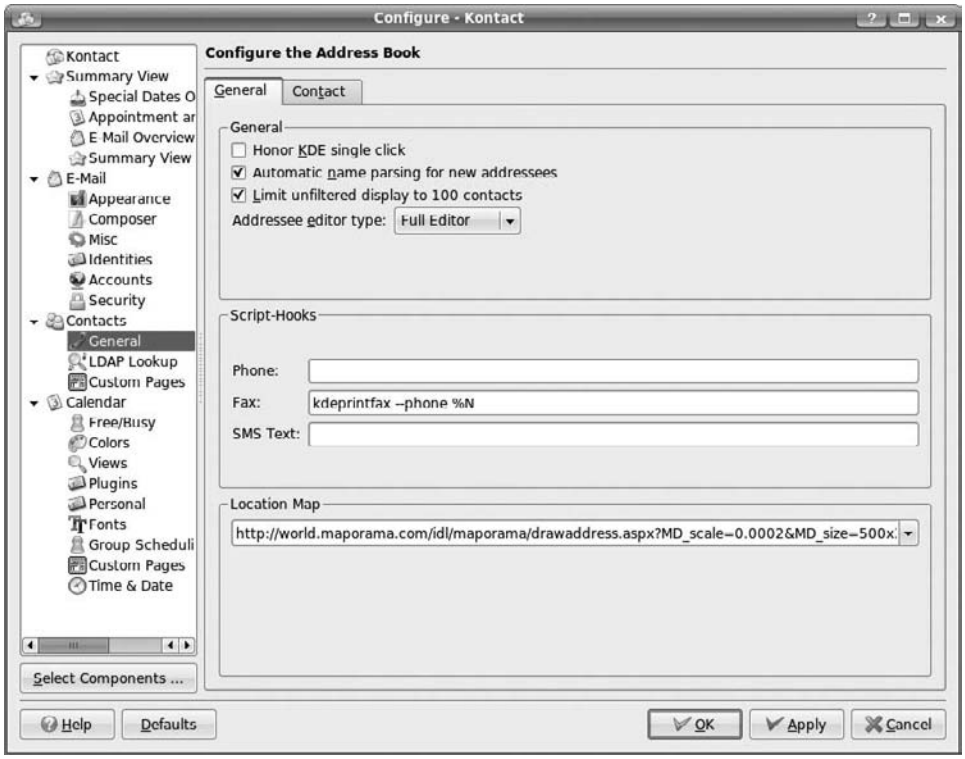
Contacts are the people with whom you often interact and exchange e-mail, and about whom you therefore want to maintain contact information: their full name; their surface address for snail mail; their e-mail address(es) at home, work, or both; their home page; any Instant Messaging addresses they have; their corporate or academic affiliation; their work, home, and mobile phone numbers; and so on. Maintaining complete and easily accessed contact information is an important part of any PIM application.

The component of Kontact that manages information about contacts is the KDE KAddressBook application, known as Contacts within Kontact. (I realize that's a bit confusing, but nobody asked me.) At any rate, the good news is that very little configuration is needed to manage Contacts within Kontact, so you won't have to endure a long section on this topic. The Contacts address book is also used by other KDE applications, such as the Kopete Instant Messaging client that is discussed in Chapter 13.

To begin configuring the Kontact address book, select the Contacts ➪ General menu item from the left pane of the configuration dialog shown in Figure 10-2. A screen like the one shown in Figure 10-14 displays, which enables you to customize certain aspects of how the Kontact address book is used.

FIGURE 10-14

The General tab for configuring contacts



To be honest, I have never had any reason to change any of the defaults for the configuration options that are displayed on this screen. (You will, however, probably want to configure the LDAP Lookup tab if you have access to an LDAP server, as explained later in this section.) Regardless of my personal configuration needs, the options on this screen enable you to single-click on contact names to edit them; decide how new names are parsed when added to the address book; define the maximum number of contacts to display by default; specify applications to use when sending appropriate messages to contact phone numbers, fax numbers, and SMS text messages; and so on.

The Contact tab enables you to configure the default lists of Prefixes, Inclusions, and Suffixes that are available when editing a contact. Prefixes are entries such as "Mr.," "Ms.," "Dr.," and so on that can appear before the first name of an address book entry. Inclusions are words like "von" that can appear between first and last names or as part of a last name, and which are near and dear to the author's heart. Suffixes are things like "Jr.," "Sr.," "III," and so on.

TIP

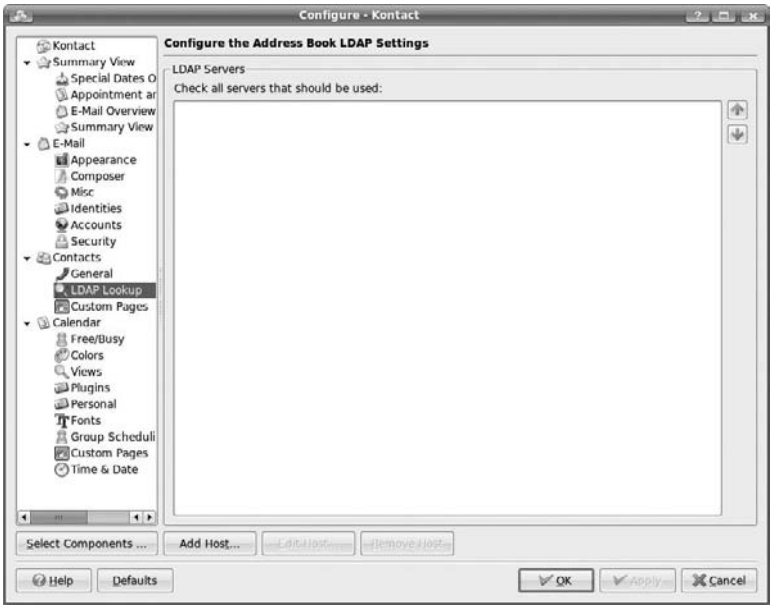
Inclusions are supported because different countries sort names that contain inclusions differently. For example, in the United States, my last name is "von Hagen," and is thus sorted under "v." In Germany, my last name would be listed as "Hagen, von," and would therefore be sorted under "H."

Companies and academic sites often store address book information in centralized, network-accessible locations. The usual way of accessing these over the network is through the Lightweight Directory Access (LDAP) protocol. Contacts makes it easy for you to define these servers so that Kontact can query them for use information when sending mail (or when looking up contact information in general).

To define one or more LDAP servers to use for address book information, select the Contacts ➪ LDAP Lookup menu item in the left pane of the screen shown in Figure 10-2. An LDAP Servers summary screen like the one shown in Figure 10-15 displays.

FIGURE 10-15

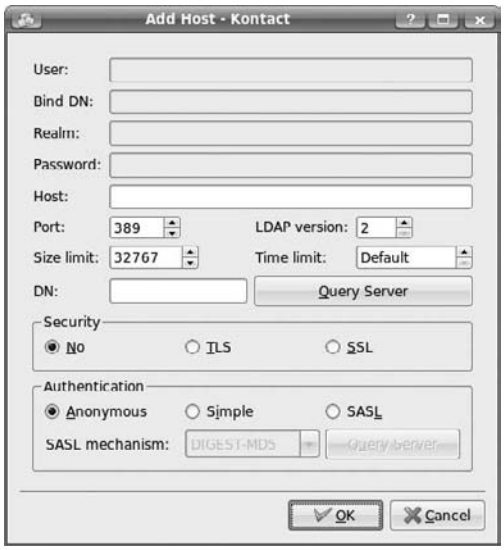
The LDAP Lookup tab for configuring Contacts



For each LDAP server that you want to define, click Add Host to display a dialog like the one shown in Figure 10-16.

FIGURE 10-16

Entering information about an LDAP server



Enter the name or IP address of the host on which the LDAP server is running, change the port that it uses (if necessary), specify the distinguished name used to locate the address book, and select any appropriate security and authentication requirements that are required in order to contact and retrieve information from the server. Click OK to accept the new LDAP server.

The name of the new LDAP server now displays in the LDAP Servers list. However, just because an LDAP server has been defined does not mean that it will actually be used—to activate a defined server so that it can be used to resolve contact/address book queries, click the checkbox to the left of its name.

TIP **When traveling, or whenever else you cannot access an LDAP server, clear this checkbox to prevent Kontact from attempting to do lookups against an unreachable server.**

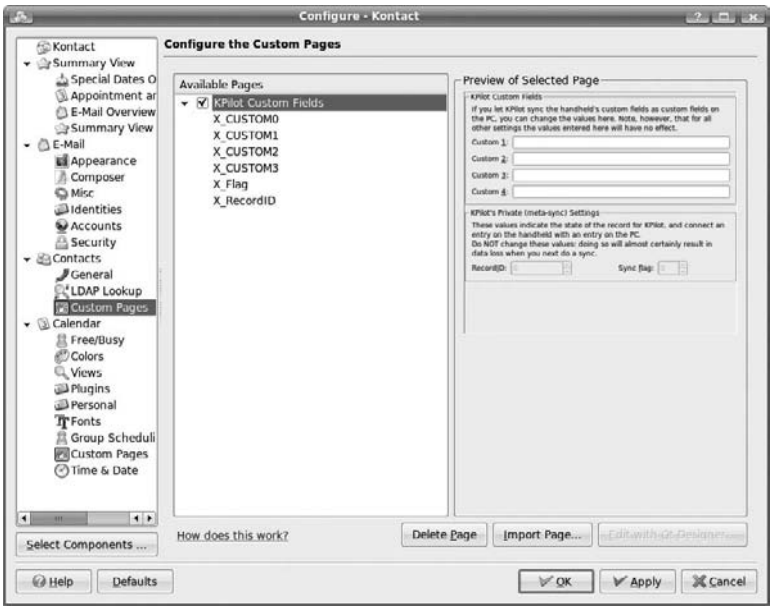
After defining all of the LDAP servers that you want to use, click Apply to accept the new values and continue making Kontact configuration changes.

The final aspect of Contacts configuration is specific to synchronizing address book data with personal digital assistants (PDAs), such as Palm handhelds. The Custom Pages tab, shown in Figure 10-17, enables you to define new fields that should be added to the contact information for each of your contacts.

Being able to add and use custom fields in Kontact enables you to use these fields to store custom fields that are specific to a given PDA application so that it will be preserved when you sync your PDA to your Contacts information, and vice versa. Because this is highly specialized, I won't go into detail here. The "How does this work?" link on the screen shown in Figure 10-17 provides a good overview of the definition process. See the KPilot home page (`http://cvs.codeyard.net/kpilot`) for more information about using custom fields.

FIGURE 10-17

Adding custom fields to contact information



Configuring Your Calendar

The component of Kontact that enables you to use and manage your calendar is the KDE KOrganizer application, known as Calendar within Kontact. Calendar also manages ToDo lists and journal entries because these are also relevant to different dates.

Very little calendar configuration is required beyond the following:

- Customizing personal information and how it is displayed
- Identifying any centralized network calendar server to which you want to publish your calendar
- Loading plug-ins for custom calendaring

All of these items are configured by selecting items in the Calendar section of the configuration dialog shown back in Figure 10-2. Many other Calendar configuration items are also available. The remainder of this section focuses on the configuration items in this bullet list because the remainder of the Calendar configuration topics are primarily related to the visual appearance and formatting of your calendar, which is both fairly straightforward and dependent on everyone's personal sense of aesthetics.

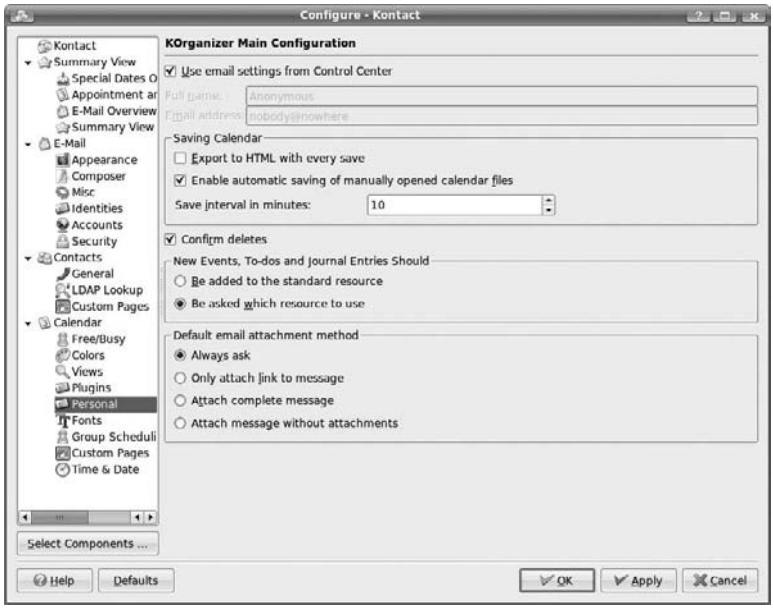
To configure the Calendar portion of Kontact, first make sure that the Calendar item in the left pane of the configuration dialog shown in Figure 10-2 is expanded.

To customize your personal interaction with the Kontact Calendar, click the Personal menu item in the left pane shown in Figure 10-2. This displays the dialog shown in Figure 10-18. The header of this dialog,

KOrganizer Main Configuration, reflects the fact that the Kontact Calendar was originally developed as a stand-alone application, but is now most easily used as part of the Kontact PIM.

FIGURE 10-18

Configuring how you interact with your calendar



The configuration items in this dialog are as follows:

- **Use e-mail settings from Control Center:** This item is a bit misleading because the KDE Control Center (which you can still execute by typing kcontrol at any command-line prompt) is being replaced by the K Menu's System Settings application. The Control Center currents gets your default name and e-mail address information from the default identity that you created when configuring e-mail in Kontact, as explained earlier in this chapter. If this item is checked, the name and e-mail address from your default identity will be used. To explicitly provide these here, clear this checkbox and supply the Full name and Email address values that you want to use.
- **Saving Calendar:** Enables you to configure how and how often your calendar is saved. Selecting the export to HTML with every save option enables you to always write out an HTML version of your calendar that you can easily include into any Web site. The "Enable automatic saving of manually opened calendar files" item enables you to specify the interval at which calendar entries that you are editing are automatically saved.
- **Confirm deletes:** Enables you to require confirmation that you actually want to delete calendar items whenever you attempt to do so.
- **New Events, To-dos and Journal Entries Should:** Kontact Calendar resources are listed in the main Calendar window. Selecting this option causes you to be prompted for which resource to

save your changes to. (You will not be prompted for a resource if only one resource is defined, which is the default behavior for the Kontact Calendar.) For more information about resources, see the section "Using and Managing Your Calendar," later in this chapter.

- **Default email attachment method:** Enables you to configure how Calendar items will be delivered to others who are scheduled for various events.

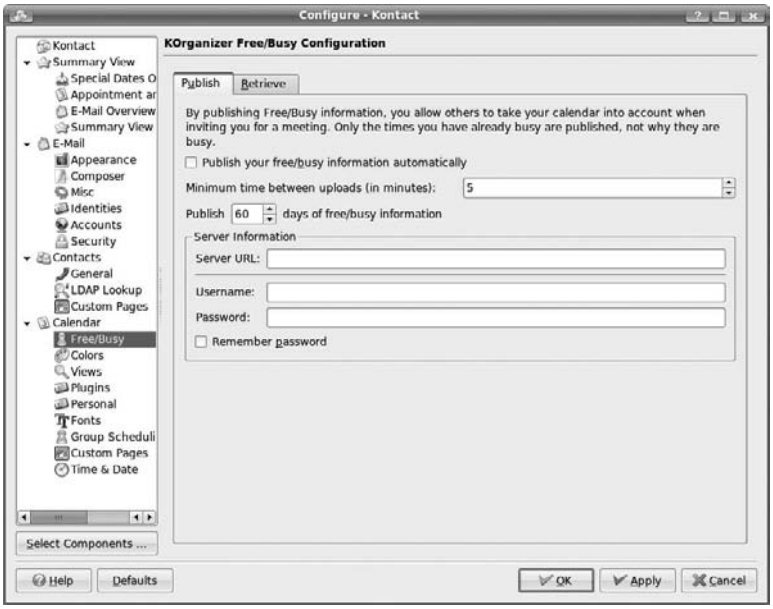
After making any changes to this screen, click Apply to save your choices and continue configuring the Kontact Calendar, or click OK to save your changes and return to the Kontact Summary View.

In modern networked environments, one of the key features of electronic calendaring is the ability to publish your personal calendar so that others can access it to find out when they can contact you, set up meetings where your presence is desired, and so on. This is done by publishing your calendar so that it is accessible to others, and by marking portions of that calendar as Free/Busy time. If your calendar is up-to-date, free time identifies times at which you are not currently scheduled for anything. Busy time is the reverse — times at which you are already committed.

To configure how your free/busy time is published (and where it is published to), and to configure the location(s) from which you can retrieve other people's free/busy time, select the Free/Busy menu item in the Calendar section of the Configure Kontact dialog shown in Figure 10-2. A dialog like the one shown in Figure 10-19 displays.

FIGURE 10-19

Defining a network calendar for posting Free/Busy information



TIP **Using the dialog shown in Figure 10-19 to configure Free/Busy time is necessary in shared, network calendaring environments only. If your site does not run a calendaring server of some sort, use the Personal configuration screen's Saving Calendar option to automatically save an HTML file that you can push to a Web server, or modify/define a KDE resource to save your calendar data in a calendaring format that others can read.**

By default, the Publish tab is selected when you first display the dialog shown in Figure 10-19. This tab enables you to set the options used when writing information about your Free/Busy time to a calendaring server. The options in this dialog are as follows:

- **Publish your free/busy time automatically:** Enables you to specify that your calendar information will be published to a specified URL automatically.
- **Minimum time between uploads (in minutes):** Enables you to specify the minimum amount of time that can pass before your calendar information is re-uploaded. This prevents you from wasting computer time and cycles by updating your calendar too frequently when you are actively updating it.
- **Publish *n* days of free/busy information:** The number of days of calendar information that should be pushed to the remote server.
- **Server Information:** Enables you to specify the URL of the remote server (usually something like `http://www.vonhagen.org/groupdav` or `webdavs://www.vonhagen.org/freebusy/vonhagen@vonhagen.org`), and the Username and Password required to contact that server. You can also check the Remember password box to cause the Kontact Calendar to remember your password so that you do not have to type it each time you publish information.

To set the options used when retrieving other people's free/busy time from a calendaring server, select the Retrieve tab. This displays the dialog shown in Figure 10-20. The options in this dialog are essentially the same as those on the Publish tab, so I won't re-explain them here. Typically, the same Server, Username, and Password Information is provided in both dialogs, although this information is not reused in case they differ.

After making any changes to these screens, click Apply to save your choices and continue configuring the Kontact Calendar, or click OK to save your changes and return to the Kontact Summary View.

The last common Calendar-related configuration item is to determine whether any of the plug-ins that are available may be of use to you. To see and optionally activate the Calendar plugins, select the Calendar menu's Plugins item from the left pane of the screen shown in Figure 10-2. A screen like the one shown in Figure 10-21 displays.

The default Calendar plugins provided with Kontact are as follows:

- **Date Numbers Plugin for Calendars:** Activating this plug-in adds the sequential day number at the top of your calendar.
- **Jewish Calendar Plugin:** Activating this plug-in shows all dates in your calendar according to both the standard western and Jewish calendars.
- **Journal Print Style:** Activating this plug-in enables you to print journal entries.
- **Microsoft Exchange 2000 Plugin:** Activating this plug-in enables you to publish/retrieve calendaring information from a Microsoft Exchange 2000 groupware server.
- **Yearly Print Style:** Activating this plug-in enables you to print out a calendar for an entire year.

FIGURE 10-20

Configuring Calendar Retrieval Options

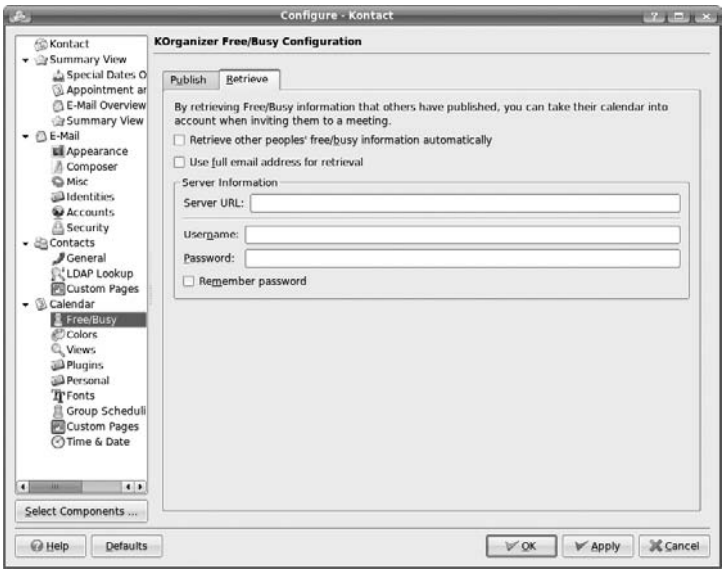
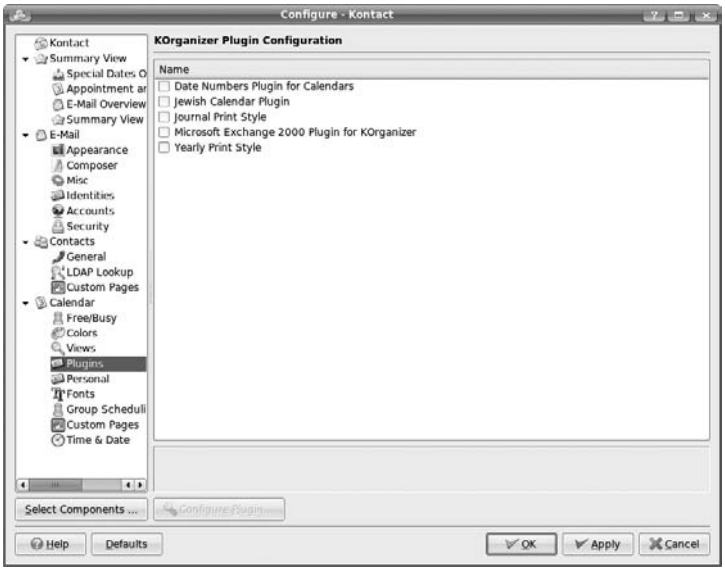


FIGURE 10-21

Configuring available plugins for the Calendar



After making any changes to this screen, click Apply to save your choices and continue configuring the Kontact Calendar, or click OK to save your changes and return to the Kontact Summary View.

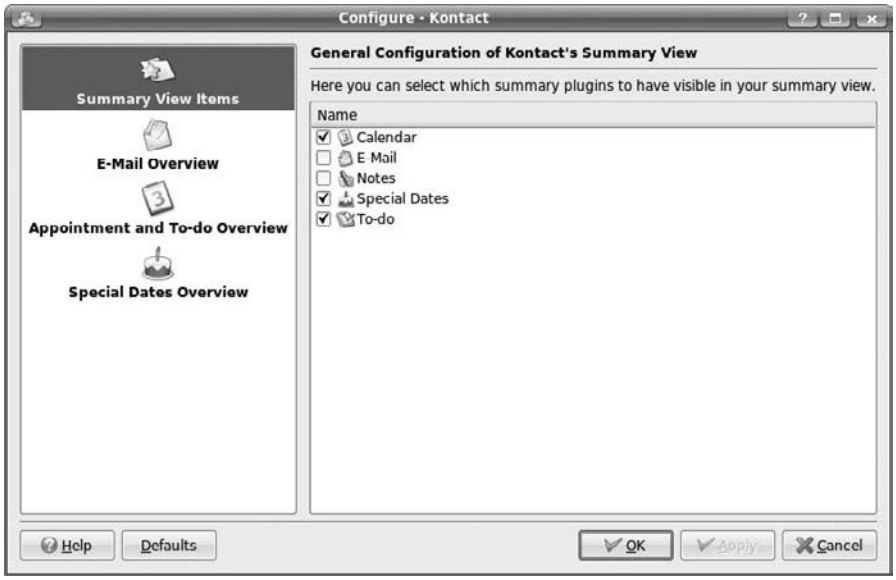
This completes the section on Calendar configuration. The next section explains how to customize the Summary View dialog that you see when you first start Kontact.

Configuring the Summary View

As you might expect from its name, Kontact's Summary View gives you an overview of new or relevant items from each of the component applications that make up Kontact. To configure the Summary View, you can use the more extensive Kontact Configuration dialog used in previous sections, or you can simply select Configure Summary View from Kontact's Settings menu. A dialog like the one shown in Figure 10-22 displays.

FIGURE 10-22

Configuring Kontact's Summary View



Each of the icons shown in the left pane of the Configure Summary View dialog enables you to customize the information that is delivered in the summary view dialog and the period of time that it encompasses. Rather than providing fairly repetitive screen shots, I'll just summarize what each of these icons enables you to configure:

- **Summary View Items:** Highlighted in Figure 10-22, this panel enables you to determine which Kontact components are listed in the Summary View. Each Kontact component includes a plug-in that delivers a configurable set of summary information.

- **E-Mail Overview:** Enables you to define which e-mail accounts, and which folders from those accounts, are listed in the summary view.
- **Appointment and To-do Overview:** Enables you to specify the time period for which appointments will be shown, and to configure whether all scheduled To-do items, or only those for the current day, will be shown.
- **Special Dates Overview:** Enables you to specify the time period for which special dates will be shown. Also enables you to configure the type(s) of special dates that will be displayed: Birthdays, Anniversaries, Holidays, and Special Occasions. This pane also enables you to configure whether special dates that are defined in any of your contacts will also be included. (The last is very handy when schmoozing business contacts.)

After making any changes to any of these screens, click Apply to save your choices and continue configuring the Kontact Calendar, or click OK to save your changes and return to the Kontact Summary View.

Sending, Receiving, and Storing Mail

Once you've configured Kontact to know about your e-mail accounts, addresses, and identities, sending and receiving mail using Kontact is as easy as you'd hope that it would be. The next few sections explain how to check for and read your mail in Kontact, how to configure aspects of Kontact associated with viewing e-mail, how to send mail in Kontact (including using the spelling corrector), and how to configure aspects of Kontact related to composing and sending e-mail.

Checking and Reading Mail in Kontact

The primary ways to check your e-mail using Kontact are as follows:

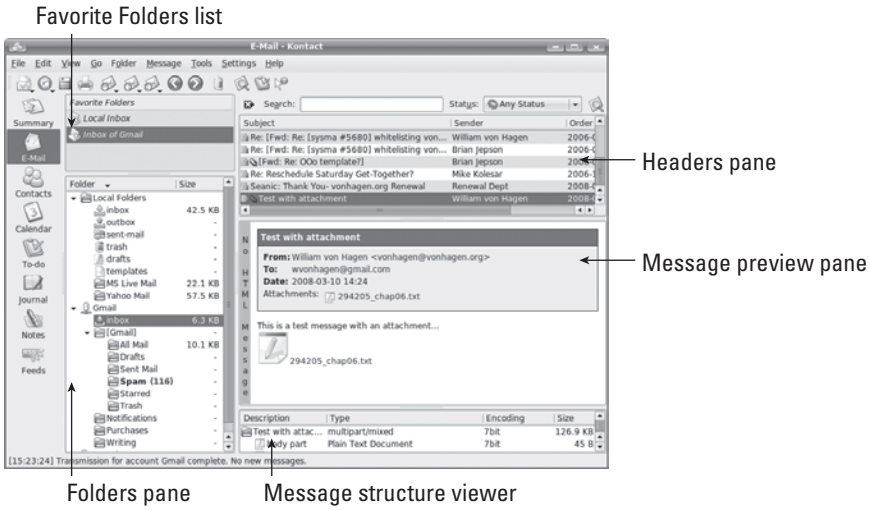
- Clicking the Synchronize All icon in the KDE 3 version of Kontact, which synchronizes the data associated with all of the Kontact components — e-mail, calendar, contacts, and so on. In the KDE 4 version of Kontact, this is referred to as the Check Mail icon.
- Clicking and holding the down arrow in the lower-right-hand corner of the Synchronize All icon within Kontact or its e-mail component. (This is the second icon from the left in the toolbar shown in Figure 10-23 — in the KDE 4 version of Kontact, this is the Check Mail icon in the toolbar.) This displays a drop-down list of individual synchronization commands. You can select Synchronize Mail (or Check Mail) from this list to synchronize only your mail (i.e., to download new messages or mail headers, and send any pending mail).
- Selecting the File ➪ Check Mail menu command (or pressing Ctrl+L) while in Kontact's e-mail component. This checks mail for all e-mail accounts that you have defined.
- Selecting the File ➪ Check Mail In command while in Kontact's e-mail component, and selecting the name of a specific e-mail account that you want to check.

TIP **To automatically check mail at scheduled internals, you must modify an account used for receiving mail, activate the Enable interval mail checking option, and specify a Check Interval. If you do not set these options, the e-mail address associated with an account will be checked only when you first start Kontact, or when you execute any of the commands in the previous list.**

Figure 10-23 shows the different panes that make up the default Kontact e-mail screen.

FIGURE 10-23

The different portions of the Kontact e-mail screen



The following list provides additional information about these panes:

- **Kontact icons:** Provide quick access to the separate applications and logical components that make up the Kontact PIM.
- **Favorite Folders list:** A list of the primary folder in which incoming e-mail is stored for each account with a distinct folder hierarchy. Multiple POP3 accounts can share the same folder hierarchy. Once a local folder hierarchy is created for a POP3 account, any folders that are subsequently associated with incoming mail from other accounts must be manually added to the Favorite Folders pane by right-clicking on the background of this pane, selecting Add Favorite Folder, and selecting the folder that you want to add from the navigation dialog that displays.
- **Folders pane:** A list of the mail folders that are available for the Kontact mail accounts on your system. This typically consists of a single set of folders that are shared by all POP3 accounts, plus one set of folders for each IMAP account.
- **Headers pane:** A list of the headers of the mail messages that are located in the folder that is currently selected.
- **Message Preview pane:** The default headers and visible portion of the body of any mail message whose header is currently selected in the Headers pane.
- **Message Structure Viewer pane:** Displays an overview of the contents of the message that is currently selected if that message is not a plain text message. (This usually means that the message is an HTML message or, more generally, a MIME message with multiple parts.) This pane is therefore not always displayed. This pane summarizes the different portions of non–plain text e-mail messages to make it easy for you to extract or examine a specific portion of the message.

As you check your e-mail, new messages are added to the message folder that is associated with each account, the name of that folder is highlighted, and the number of new messages is displayed after the name of that folder. The name of the folder where new mail for a given account is stored (usually Inbox) is listed in the Favorite Folders section of the Kontact mail screen. You can get to the Inbox folder for any account with an independent folder structure by clicking on its name in the Favorite Folders list. If you have multiple accounts for receiving e-mail, you may see multiple sets of folders listed below the Folder list.

After you select the Inbox that contains the mail for an e-mail account, the Header pane window displays the headers for any mail in that mailbox. The subject and sender of unread mail messages are displayed in bold in this pane to make it easy for you to identify new messages. By default, messages are displayed chronologically by order of arrival, with the latest mail messages shown at the bottom of the Header pane.

To view any mail message, you can either scroll through the message as displayed in the Message Preview pane, or you can double-click on the entry for that message in the Headers pane to display that message in a new window. HTML messages shown in the Message Preview pane are identified by the phrase HTML Message displayed on a vertical black background to the left of the mail header in the Message pane. Non-HTML messages are identified by the phrase No HTML Message displayed on a vertical gray background.

To reply to the current e-mail message, you can either select the Reply To icon in the toolbar (the sixth icon from the left, by default), or right-click in the message and select Reply from the pop-up menu. To reply to the sender and all of the recipients of a mail message, click and hold the Reply To icon or right-click in the mail message. You can then select Reply To All from the menu that displays.

Similarly, to forward an e-mail message to someone else, you can either select the Forward icon in the toolbar (the seventh icon from the left, by default), or right-click in the message and select Forward from the pop-up menu (or Forward To, depending on where you clicked). If you select Forward from the menu, you can choose whether you want to forward the message as an attachment, inline in the body of the new message, or as a redirect (simply re-sending the original message). If you are redirecting the mail message, a pop-up displays, prompting you for the e-mail address to which you want to redirect the message. If you are forwarding the message as an attachment or inline, a composer window displays, as discussed later in this chapter in the section "Sending Mail in Kontact." That same section also provides detailed information about composing and sending e-mail in Kontact.

To move to the next mail message in your Inbox, you can either simply select its header using the mouse, or use the left and right arrow icons in the toolbar to move to the previous or next message, respectively.

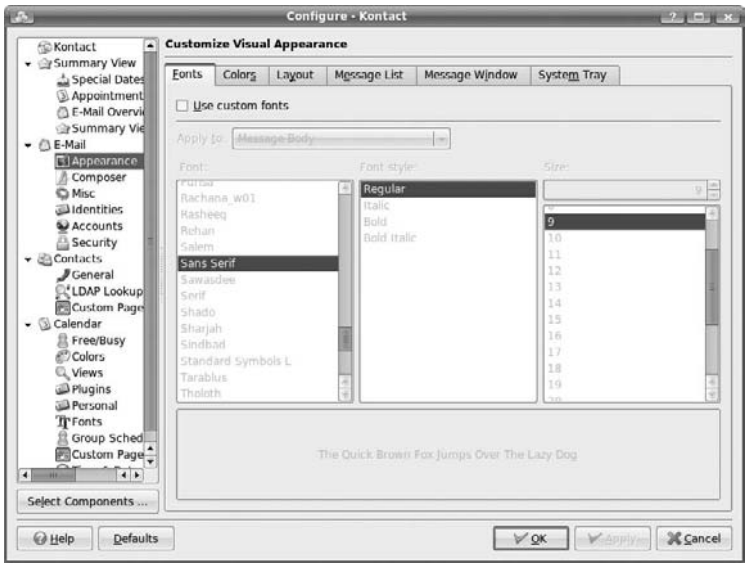
To delete any mail message, you can simply highlight the header for that message and press the Delete key on your keyboard to move that message to the Trash folder that is associated with that account. if you're a menu fan, you can also right-click on the mail header and select the Move to Trash menu command. Unless you've configured your account to automatically delete messages when you exit Kontact, you will need to manually empty the trash every so often by selecting the File ➪ Empty All Trash Folders menu command.

Setting Preferences for Reading Mail in Kontact

Kontact provides many configuration options for displaying e-mail messages and for customizing the layout of the panes that make up Kontact's e-mail component window. These can all be configured by selecting the Settings ➪ Configure Kontact menu item, and selecting the E-Mail ➪ Appearance configuration item. This displays the dialog shown in Figure 10-24.

FIGURE 10-24

Kontact's e-mail appearance configuration dialog



This dialog provides multiple tabs that enable you to configure almost everything about how Kontact's e-mail component displays messages. Many of the configuration items on these tabs are self-explanatory (such as the font configuration options shown on the default tab). Rather than waste your time and bloat this book by explaining each and every option, I'm going to focus on the ones that I often change or at least think are the most interesting. Your mileage may vary (but hopefully not too much).

If you have many mail folders (as I do), even the subset of those folders that is displayed in the Folders pane isn't enough to see all of them. I generally prefer the ability to see a larger preview of my mail messages as opposed to a long list of folders. Similarly, I don't find the Messages Structure View pane to be all that useful (although it is interesting). Apparently, many other Kontact users agree with me because these aspects of the Kontact display are simple configuration items in Kontact.

To configure some common aspects of the panes that are displayed in Kontact from the Appearance dialog shown in Figure 10-24, do the following:

1. Click the Layout tab. The configuration dialog shown in Figure 10-25 displays.
2. To show a shorter list of folders, a smaller Folder pane, and a larger Message Preview pane, click the Short folder list item in the Folder list section in this dialog.
3. To eliminate the Message Structure Viewer, click Show Never in the Message Structure Viewer section.

After making these selections, click Apply to accept the changes and continue making configuration changes to Kontact's e-mail component, or click OK to accept your changes and close the configuration dialog. After making these changes, your Kontact e-mail component will look like the window shown in Figure 10-26.

FIGURE 10-25

The Layout tab in Kontact's e-mail appearance dialog

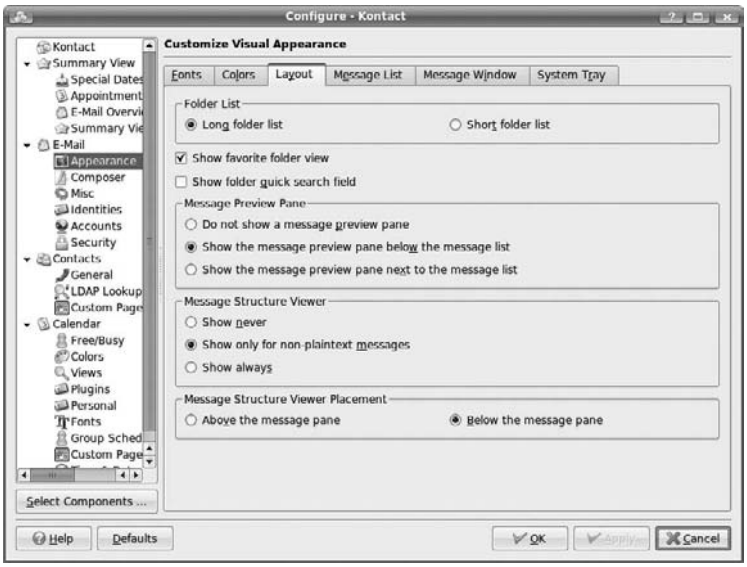
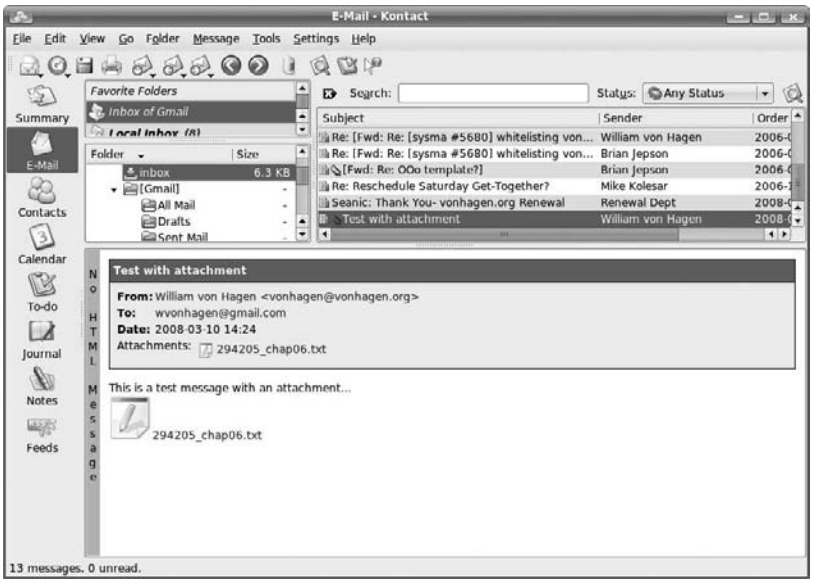


FIGURE 10-26

Kontact's e-mail component after configuration changes



The sample message shown in this figure doesn't really benefit from the additional message preview space, but you can certainly see that long messages from your friends, mailing list digests, and your favorite comic strips certainly would.

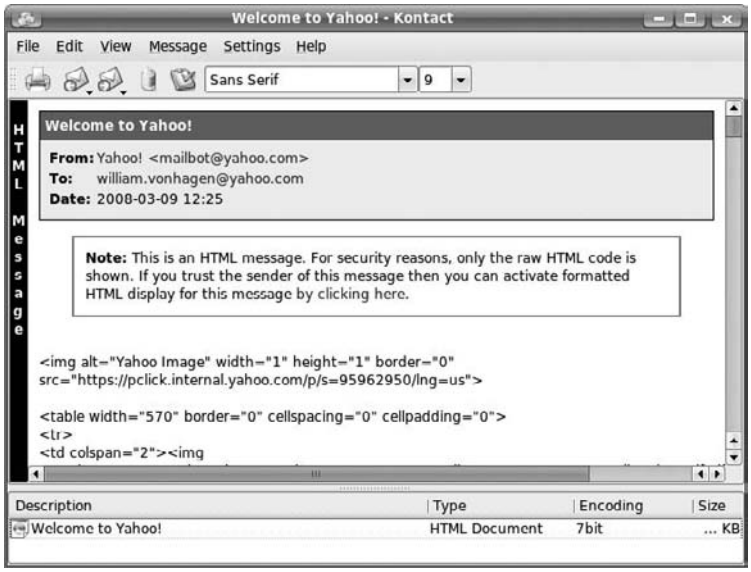
TIP **If you have only a single e-mail account, the Favorite Folders pane may not be all that useful to you. To eliminate this, clear the Show favorite folder view checkbox shown in Figure 10-25 and click Apply or OK.**

If you participate in e-mail exchanges with lots of back and forth between recipients, you may want to be able to group those messages together. This is called threading, which means that each "thread" of conversation (in this case e-mail) is organized hierarchically under the first message in the thread. Threading makes it very easy to follow e-mail conversations because each message is displayed indented below the message that it was in response to. This is especially handy when multiple people respond to different e-mail messages on a topic—using threading makes it easy to see this fragmentation and figure out who was responding to what. To activate threading in Kontact, select the Threaded Message List option on the Appearance ➪ Message List tab, and click Apply or OK. The messages in your Inbox (or other folders) will automatically be reorganized to show the message threads. If you decide that you don't like threading, you can always clear this option and click OK or Apply, and all of your messages will be displayed using the default chronological order.

The final item on my configuration hit list is the fact that, for me, one of the more disappointing things about reading mail in Kontact is that HTML messages are initially displayed as shown in Figure 10-27.

FIGURE 10-27

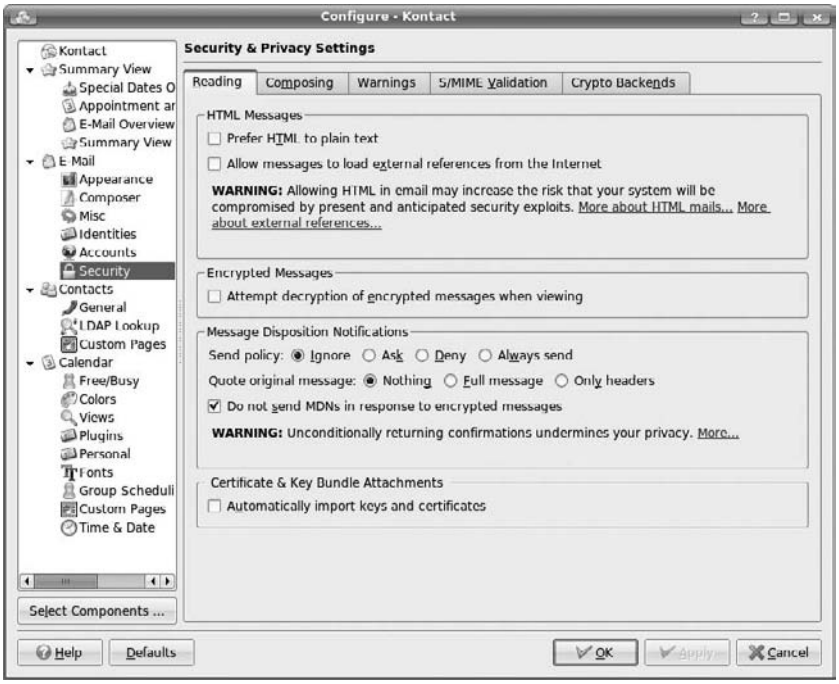
Kontact's default display of HTML e-mail



I don't send HTML e-mail, but many people insist on doing so, as do news sites that keep me up-to-date with the troubles of the world. While it's nice that Kontact is trying to protect you from nasty embedded HTML code, this hand-holding is more often irritating than it is useful. This is a Linux system, after all, not some vulnerable and pathetic Windows box. You can always display a single message using its HTML formatting by first clicking on the link in the warning to display HTML formatting, and then by clicking the link in the second warning to load external references, but that can be quite time-consuming if you get a lot of HTML mail. Therefore, I usually enable the automatic display of formatted HTML messages. To do this, first select the Settings ➪ Configure Kontact menu item, and select the E-Mail ➪ Security configuration item. This displays the dialog shown in Figure 10-28.

FIGURE 10-28

Kontact's E-Mail Security configuration dialog



To automatically display HTML messages using the formatting that they contain, select the Prefer HTML to plain text checkbox and click OK. HTML mail messages will then look something like the window shown in Figure 10-29.

This is an improvement over the same message as shown in Figure 10-27, but you can see that a warning message is still being displayed. To eliminate this message and simply show the HTML message as the sender intended, you must also click the "Allow messages to load external references from the Internet" checkbox shown in Figure 10-28. Doing so causes HTML e-mail to display without any error messages, as the sender intended, as shown in Figure 10-30.

FIGURE 10-29

Formatted HTML e-mail in Kontact



FIGURE 10-30

Formatted HTML e-mail with external references in Kontact



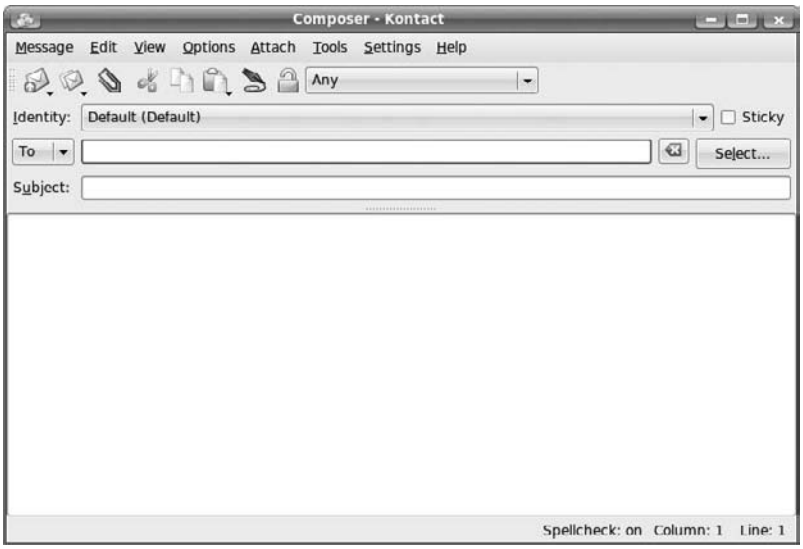
CAUTION **As the warning shown in Figure 10-28 implies, allowing HTML messages to resolve and load images and other external references in HTML e-mail can have side effects. This can enable HTML messages to load malicious code to exploit display and rendering problems, load bogus links that you could accidentally click on, and might also enable the people who send HTML e-mail to know that you have read it. For more information, click on the warning links shown in Figure 10-28. I find that the convenience of actually being able to read news summaries outweighs any worries I may have about those warnings, but you may be more concerned about privacy and security. This is a personal, individual decision. Make sure that you read the warning messages shown by clicking on the links in Figure 10-28 and understand the risks before permanently enabling HTML display and external reference loading. Caveat emptor!**

Sending Mail in Kontact

To send a new mail message, make sure that Kontact's e-mail component is displayed by clicking the Mail icon (second from the top) in the left pane of the summary view or any Kontact component. Click the New Mail icon (the left-most icon in the Kontact E-Mail component window). If you're a menu fan, have disabled the toolbar in Kontact's E-Mail component, or are in Kontact's Summary View, you can also begin sending a message by selecting the File ➪ New Message menu item. Regardless of which option you select, a mail composition window like the one shown in Figure 10-31 displays.

FIGURE 10-31

Sending mail using Kontact



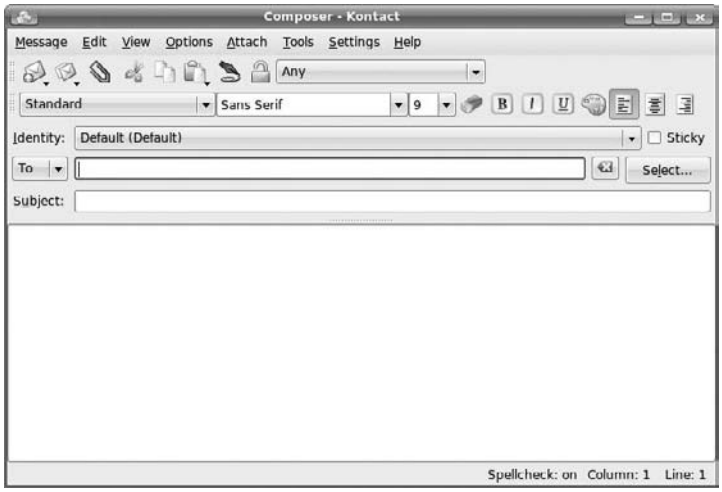
TIP **To reply to or forward an existing e-mail message, click the Reply To or Forward icons in the toolbar (the sixth and seventh icons from the left, by default), or right-click in the message and select the Reply or Forward menu items from the pop-up menu. Both of these menu items display sub-items that enable you to specify whom you want to reply to (the sender, all recipients, a mailing list, and so on) or how you want to forward the message, respectively.**

Enter the e-mail address to which you're sending the message in the To field. Kontact will attempt to auto-complete the e-mail address using your contacts, as well as an internal list of the addresses to which you have recently sent mail. If you see the e-mail address that you want to send mail to in the completions drop-down list, you can select that address at any time by left-clicking on it. You can also click Select… to display a dialog that enables you to browse the contents of your address book and select the people to whom you want to send your mail.

Next, enter the subject and text of your message. If you insist on sending HTML mail, you can select the Options ➪ Formatting (HTML) menu command to display a formatting toolbar that contains drop-down style and font menus and some basic formatting icons, as shown in Figure 10-32.

FIGURE 10-32

The HTML Formatting toolbar for composing HTML e-mail



TIP **After you enter a subject for your message, that subject will replace the word Composer in the mail composition dialog header—this helps you differentiate among multiple composition windows if you happen to be working on more than one message at the same time.**

Automatic spell-checking is enabled by default in Kontact, so any words that are not found in the default dictionary will be highlighted in red. To see suggested corrections for any of these words, right-click on the word to display a pop-up list of suggested corrections. If the word is indeed misspelled and a correction is available, simply select it from the list. If the word is not misspelled and you want to add it to your dictionary, select the Tools ➪ Spelling menu item to display the dialog shown in Figure 10-33.

Click Add to Dictionary to add the word to your personal dictionary and move on to any subsequent misspelled words. You can also click Replace or Replace All to correct this term once or all times that it appears in your message. Similarly, you can select Ignore or Ignore All to tell the spell-checker to allow one or all instances of the potentially misspelled term in your message. The latter can be useful if, for example, your mail contains technical, variable, or other specialized references that you don't want to permanently add to a dictionary.

FIGURE 10-33

Adding words to the Kontact dictionary



If no suggested correction is accurate, type an accurate correction in the Replace with: text entry box. You can then click Replace or Replace All to replace one or all instances of the term with the accurate spelling.

To attach a file to your mail (we're back in Figure 10-32 now), click the paperclip icon. This displays a browse dialog that makes it easy to navigate to and select the file that you want to attach. You can also open a file manager such as Dolphin or Konqueror and drag and drop the file into your mail message, but I find this time-consuming because I don't usually have a file manager running.

To send your mail as the e-mail user associated with your default identity, click the Send Mail Via icon (the left-most icon) or select the Message ➪ Send Mail menu command. To send the mail using another identity, click and hold the Send Mail Via icon and select the identity that you want to send mail as from the drop-down list. You can also select the Messages ➪ Send Mail Via menu command and select the identity that you want to use from the drop-down list that displays.

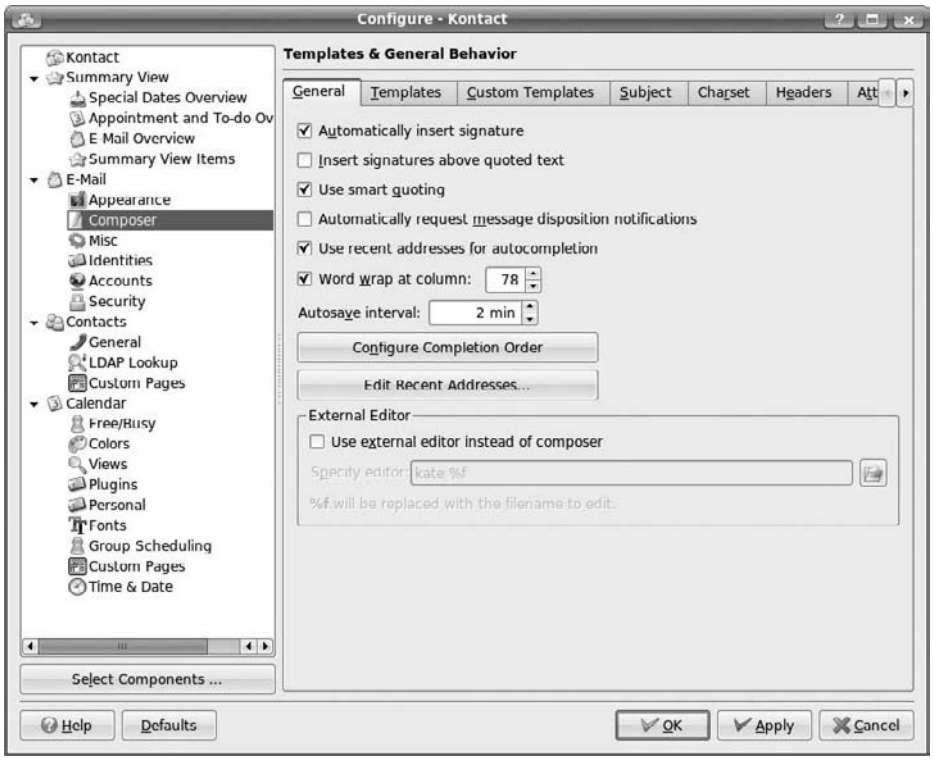
If you are working offline or simply don't want to send mail immediately, you can click the Send Mail Later icon (the second icon from the left) or select the Message ➪ Send Mail Later menu command. These queue the messages for sending. (To send the mail later using an identity other than your default identity, click and hold the Send Mail Later icon or the Message ➪ Send Mail Later Via menu command, and select the appropriate identity from the drop-down list.) You can subsequently send these queued messages using Kontact's File ➪ Send Queued Messages menu command.

Setting Preferences for Sending Mail in Kontact

While I usually customize some aspects of reading and viewing mail in Kontact, I'm pretty happy with the default values that Kontact uses when composing and sending mail. However, like everything about e-mail in Kontact, you can customize many aspects of composing and sending mail by selecting Settings ➪ Customize Kontact menu item and selecting the E-Mail ➪ Composer item in the left pane. (If you are already in the E-Mail application, you can also select Settings ➪ Customize KMail and click the Composer icon.) This displays the dialog shown in Figure 10-34.

FIGURE 10-34

Configuring mail composition in Kontact



TIP **You can also configure security settings related to composing mail by selecting the Security item or icon and clicking the Composer tab.**

Rather than going through all of the possible configuration options for Kontact's e-mail application and losing you in a sea of things that you probably won't want to change, I'll just discuss my favorite. You may be able to find many others by exploring the various tabs of Kontact's Composer configuration item.

I've been using various flavors of Emacs for 25 years now, and my fingers just work that way. To me, Ctrl+A will always mean "Go to the beginning of the current line," and selecting an entire mail message a thousand times by accident hasn't cured me. Call me old-fashioned, but I'd prefer to use the keyboard combinations that I already know whenever possible, especially because I don't send HTML e-mail, and therefore don't need any special toolbar and menus to simplify things.

For people like me, Kontact's e-mail application thoughtfully enables you to customize the editor that you use when composing e-mail. To do this, first display the Composer customization window shown in Figure 10-34. By default, the General tab displays. Next, click the Use external editor instead of composer checkbox. The Specify editor text entry box activates, enabling you to specify the editor and command that you want to use. By default, this area displays `kate %f`, where `%f` represents the name of the tem-

porary file where Kontact will store your e-mail as you write it. In my case, I change this to `emacs %f`, but you're welcome to use any editor that you feel comfortable with.

After making this setting, the Composer window shown earlier in Figure 10-31 will display whenever you begin to compose a mail message, so that you can enter the To and Subject fields normally. However, positioning the cursor in the message body window and typing any standard character will automatically start the editor that you've specified. After composing your message in that editor, simply exit the editor normally, and whatever you've typed will be inserted into the Composer's message body pane. You can then send your message normally.

TIP **If you're even more of a purist than I am and want to use the vi editor to compose your messages, you'll have to use a command such as `konsole -e vim -f %f` as your editor command because vim is a command-line application. For a graphical vi-like editor, you can install gvim (via the vim-full or vim-gnome packages, which will also install lots of GNOME dependencies), try to find a copy of the aging Kvim, or build and install a modern vi-like editor such as Yzis (`www.yzis.org`).**

Organizing and Filtering Mail

The next few sections provide suggestions for organizing your e-mail by using folders, setting up e-mail filters to automatically file incoming e-mail into different folders, and configuring Kontact to catch and automatically manage incoming mail that can be identified as junk mail, better known as spam, or as containing malicious attachments such as viruses. Both identifying spam and identifying mail that may contain a virus require the use of external tools, which can easily be integrated into Kontact through the use of built-in wizards.

Using Folders to Organize Your Mail

As more and more communication becomes electronic, it becomes increasingly useful, often necessary, to keep some of your e-mail around long after it was actually sent. For example, many online purchase receipts are e-mail-only, and I haven't received paper mail from friends and family members in years. Keeping all of this in your Inbox is possible, thanks to Kontact's e-mail search capabilities; and display options such as threading (discussed in the section "Setting Preferences for Reading Mail in Kontact") can simplify finding related messages. However, the best solution to grouping and locating related messages is to create folders within Kontact and use those to store related messages. This is conceptually identical to using folders/directories to organize the files (and directories) in your home directory. I use folders such as Personal to store personal e-mail messages, Purchases to store e-mail purchase receipts, and so on. I even have folders within folders so that I can better organize correspondence with different friends, family members, and so on.

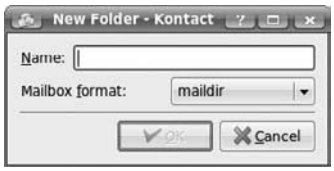
To create a folder, right-click the existing folder in which you want to create the folder. If you are using POP3 mail, this is probably the Local Folders folder (at least at first; you can subsequently create folders within other folders if you want to). If you are using IMAP mail, this is probably the name associated with your incoming mail account.

Select the New Subfolder menu item. A dialog like the one shown in Figure 10-35 displays.

Figure 10-35 shows a new local folder being created for a POP3 account, which gives you the opportunity to specify the format of that folder; if you are creating a folder for an IMAP account, no format option is displayed because folders are logical entities stored on the mail server, not local files/directories.

FIGURE 10-35

Creating a new mail folder in Kontact



To complete the creations of the folder, enter the name of the new folder and click OK.

Once you've created a folder, you can move mail messages to it by simply dragging and dropping one or more message headers from Kontact's header pane onto the new folder. By default, Kontact will display a pop-up menu that enables you to select whether you want to copy or move the message(s) to the folder.

If you find the pop-up menu irritating because you always want to move messages to a folder when dragging and dropping them, you can disable the menu by selecting Settings ➪ Configure Kontact, selecting the E-Mail ➪ Misc section, and clearing the "Ask for action after dragging messages to another folder" checkbox. After you make this change, messages will simply be moved to a folder when you drag and drop them there.

Using Filters to Process Incoming Mail

Filters enable Kontact's e-mail application to automatically analyze incoming e-mail and process it in different ways based on portions of its content. By default, Kontact enables you to create filters based on the mail's recipient (the "To" field in the mail header), from the mail's sender (the "From" field in the mail header), or the subject of the e-mail (the "Subject" field in the mail header). Kontact can also detect mail that is sent to a mailing list (to which you are presumably a subscriber), and supports filtering based on the generic "To" field for such messages.

You can create filters in one of two ways: either by right-clicking on a message that is an example of a message that you want to filter or by selecting the Message ➪ Create Filter menu command. Selecting either of these displays a secondary menu from which you can select Filter on Subject, Filter on From, Filter on To, or Filter on Mailing List to create a filter based on the associated field. After you select any of these, a Filter Rules dialog like the one shown in Figure 10-36 displays.

Figure 10-36 shows the Filter Rules dialog for a mailing list filter, which is listed in the Available Filters pane. The General tab, at right, shows the basis of the filter (the Filter Criteria) and the actions (Filter Actions) that should be taken when the filter criteria are matched. In this case, the filter criterion is that the filter is applied if a message matches all of the available conditions. In this case, these conditions are simply that the List-Id must contain the name of the mailing list. Defining multiple criteria enables you to create very granular filters that require that multiple criteria be matches, or to apply a single filter if the message matches any of a number of conditions (in the latter case, after selecting the "Match any of the following" radio button).

When you're first creating a filter, the Filter Actions portion of the Filter Rules dialog is incomplete because Kontact has no way of predicting what you want to do in response to matching the criteria. The most common use of a filter is to file messages in a specific folder if they match the filter criteria, so this action is initially suggested, although you must specify a folder by clicking the folder icon to the right of Please Select a Folder and navigating to the folder into which you want these messages to be filed.

FIGURE 10-36

Creating a filter in Kontact's e-mail application



To add another action, click More. This adds an empty filter action, which you can configure by clicking on that action and selecting an action from the drop-down list. This list includes items such as Move Into Folder, Copy Into Folder, Mark As, Forward To, Confirm Delivery, Play Sound, and so on. Selecting any of these displays another drop-down, from which you can select or specify the target of the selected action. For example, when filing mail based on a mailing list, a common action is to mark the message as being read, which you can do by selecting the Mark As action. A drop-down list of available status types displays to the right of this action, from which you can select Read.

The Advanced tab, shown in Figure 10-37, enables you to configure various aspects of how and when the filter is applied.

In general, the default settings provided on this pane are fine. In some cases, you may want to configure Kontact so that this filter is applied only to incoming mail from certain accounts, which you can do by selecting the From checked accounts only radio button and de-selecting the accounts to which you do not want to apply the filter.

Another interesting setting on the Advanced tab is the "If this filter matches, stop processing here" checkbox, which is active by default. This option causes subsequent filters to be skipped once the current filter is matched. In some cases, you may want to create a filter that does some processing on a message, with subsequent filters depending upon that processing. If you want such cascading filters to apply, you must clear this checkbox.

Once you have created a filter, click Apply to save your changes but continue editing the rule, or click OK to close the Filter Rules dialog and return to Kontact.

To edit existing filters, select the Settings ➪ Configure Filters menu option from within Kontact's e-mail application. This displays a dialog like the one shown in Figure 10-36, but in which you can select any existing rule by selecting its identifier from the Available Filters pane. Filters are applied in the order

shown in the Available Filters list. You can also use this dialog to change the order in which filters are applied by selecting the name of that filter in the Available Filters pane and using the up and down arrows or beginning/end of list buttons to move the rule up or down within the list. After changing the rule sequence or anything about any of the rules, click Apply to save your changes but continue editing your filters, or click OK to close the Filter Rules dialog and return to Kontact.

FIGURE 10-37

Advanced settings for a filter in Kontact's e-mail application



Configuring Spam and Virus Detection in Kontact

Kontact's e-mail application can be configured to automatically invoke external software that is designed to detect spam, mail containing viruses, and so on. These applications can easily be configured using the Anti-Spam Wizard and Anti-Virus Wizard commands on the Tools menu. Integration with Kontact's e-mail application is done by creating custom Kontact filter rules that run the specified application and perform the actions that you have requested on matching messages.

By default, a popular anti-spam application known as Bogofilter (`http://bogofilter.sourceforge.net`) is installed on Kubuntu systems for spam detection and processing. No anti-virus software is installed by default, although packages such as Clam AV (`www.clamav.net`), Sophos RST-B (`www.sophos.com/rst-detection-tool`), F-Secure's Anti-Virus tool (`www.f-secure.com/small_businesses/evaluations`), and so on are available for Linux. I personally run anti-spam software only on my desktop Linux systems, but you can certainly run anti-virus software on mail servers, as described in Bonus Chapter 9.

To integrate Bogofilter with Kontact's e-mail application, select the Tools ➪ Anti-Spam Wizard menu item. The wizard searches your disk for installed anti-spam software and displays a list of matching software, as shown in Figure 10-38 (which contains only Bogofilter, by default).

FIGURE 10-38

The initial dialog of Kontact's anti-spam wizard



Click Next to proceed. A second dialog displays, shown in Figure 10-39, in which you can configure the actions that should be taken when Bogofilter identifies incoming mail as spam.

These default actions, which are implemented by Kontact filter rules, mark the incoming spam as having been read and move it to the Trash folder for your account. You can change the folder to which spam is moved by clicking the folder icon in this dialog and selecting another folder in the dialog that displays. I typically create a spam folder and move messages there, to give me the opportunity to occasionally scan the folder to ensure that mail from friends of mine isn't actually being marked as spam. I have to do this because one of my friends has the unfortunate name of "Make Money Fast."

After making any changes that you want to make, click Next to accept these actions and proceed to the next dialog. This dialog is simply a confirmation dialog, which I'm not going to show because it's pretty trivial. You have three choices in this dialog:

- Return to a previous dialog by clicking Back so that you can make additional changes.
- Accept anti-spam tool integration by clicking Finish.
- Abort anti-spam tool integration by clicking Cancel.

In most cases, click Finish. You can see the new anti-spam Filter rules that have been created by selecting the Settings ➪ Configure Filters menu command. Figure 10-40 shows a Filter Rules dialog that contains Bogofilter rules.

You can compare this dialog against the one shown in Figure 10-36 to see anti-spam rules, although which ones they are is pretty clear.

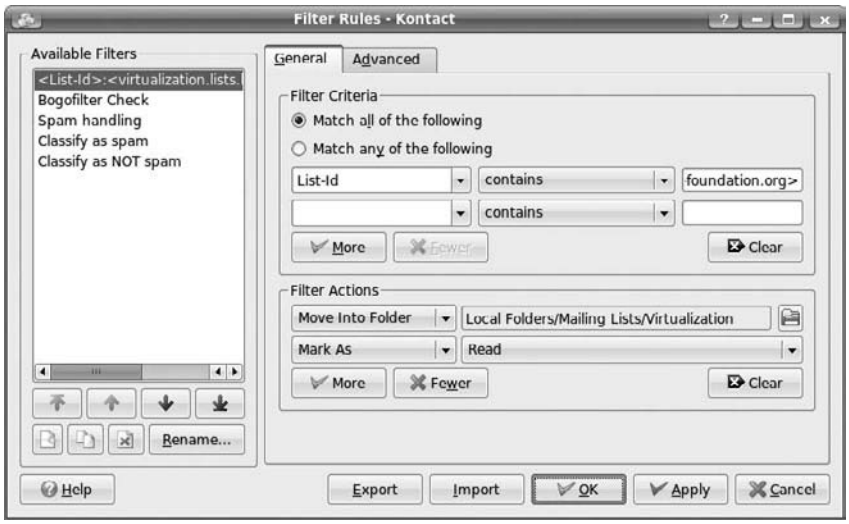
FIGURE 10-39

Specifying actions to be taken for spam



FIGURE 10-40

Bogofilter anti-spam rules in Kontact



TIP **Another popular anti-spam tool is SpamAssassin, which is commonly used with Evolution, the GNOME mailer used on standard Ubuntu systems that was discussed in Chapter 9. You can use SpamAssassin with Kontact by installing SpamAssassin using the Adept Installer, deleting any anti-spam rules that use Bogofilter, and then re-running the Anti-Spam Wizard and selecting SpamAssassin as your anti-spam tool. (You should never run two spam detection tools in Kontact at the same time because their filter rules may interact in strange ways.) The wizard will then create new, equivalent rules that use SpamAssassin for spam identification and processing. SpamAssassin uses a somewhat different model for spam identification than Bogofilter, and therefore requires training in order to work efficiently. If you are not satisfied with the performance of Bogofilter, you should definitely experiment with SpamAssassin.**

Creating and Managing Contacts

As discussed at the beginning of this chapter, maintaining information about personal and business contacts is an important part of a personal information management suite such as Kontact. The most basic information about a contact is a person's name and e-mail address, but you can also save information such as their home address, phone numbers, other e-mail addresses, and so on. Adding information about your contacts is useful not only in providing a single source of information about them, but also helps when sending e-mail to those people, because the list of available contacts provides default suggestions for e-mail address completion when you are composing e-mail messages. Contact management is supported in Kontact through an integrated application known as KAddressBook, which is simply known as the Contacts application in Kontact.

Kontact makes it easy to add basic name/e-mail information about the sender of incoming mail by enabling you to right-click on the name or e-mail address of the sender and selecting Add to Address Book from the pop-up menu shown in Figure 10-41.

Adding contact information to your address book in this fashion initially displays a dialog that confirms that the contact has been added and informs you that you can add more information about that contact from within the address book itself.

To access your address book, whether to add a new contact or edit information about an existing contact, select the Contacts icon from Kontact's icon pane. Figure 10-42 shows the default contacts screen, which lists all of the contacts that are currently defined.

To add information about a new contact, right-click in the Contacts list and select New Contact from the pop-up menu that displays. After adding information about the new contact, you can then click Apply to save your changes and continue editing, or click OK to save your changes and close the detailed contact information dialog.

To edit information for any existing contact, double-click on the name of that contact to display a screen that is similar to the one shown in Figure 10-43, but in which certain fields already have values for the contact that you selected. After adding or changing information in various fields, click Apply to save your changes and continue editing, or click OK to save your changes and close the detailed contact information dialog.

FIGURE 10-41

Adding the sender of incoming mail as a contact

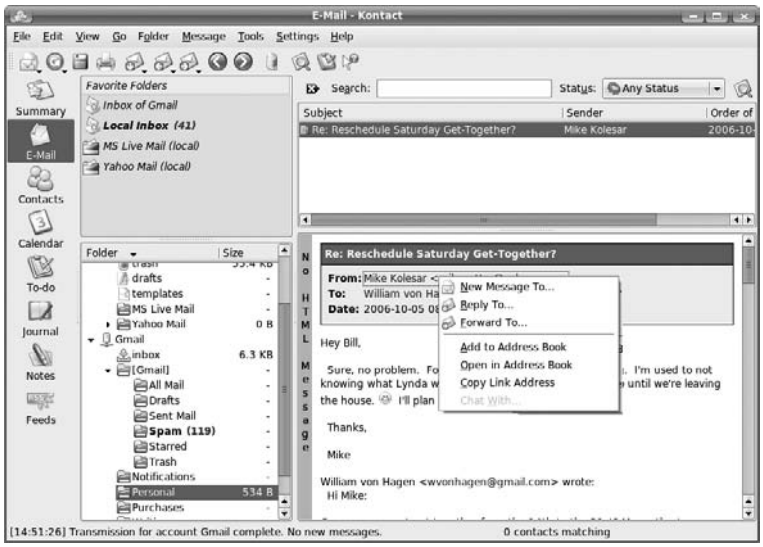


FIGURE 10-42

The Contacts/Address book screen in Kontact

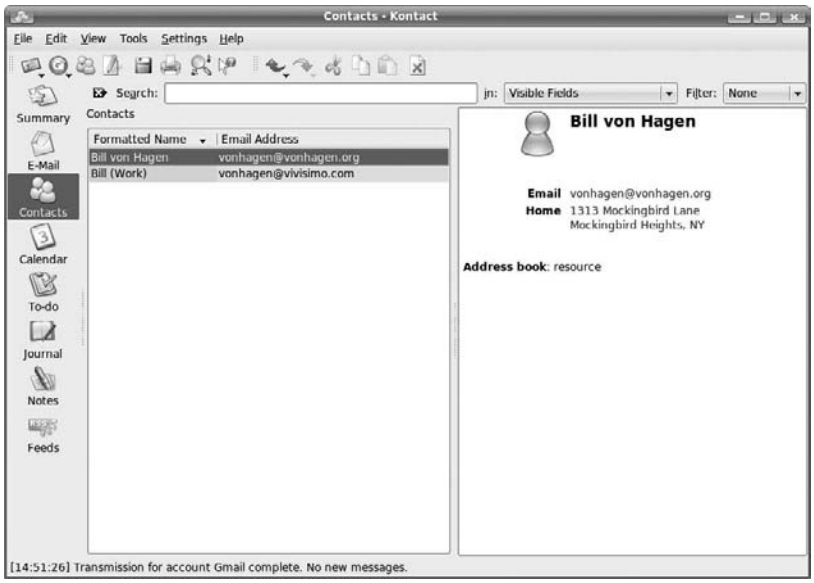
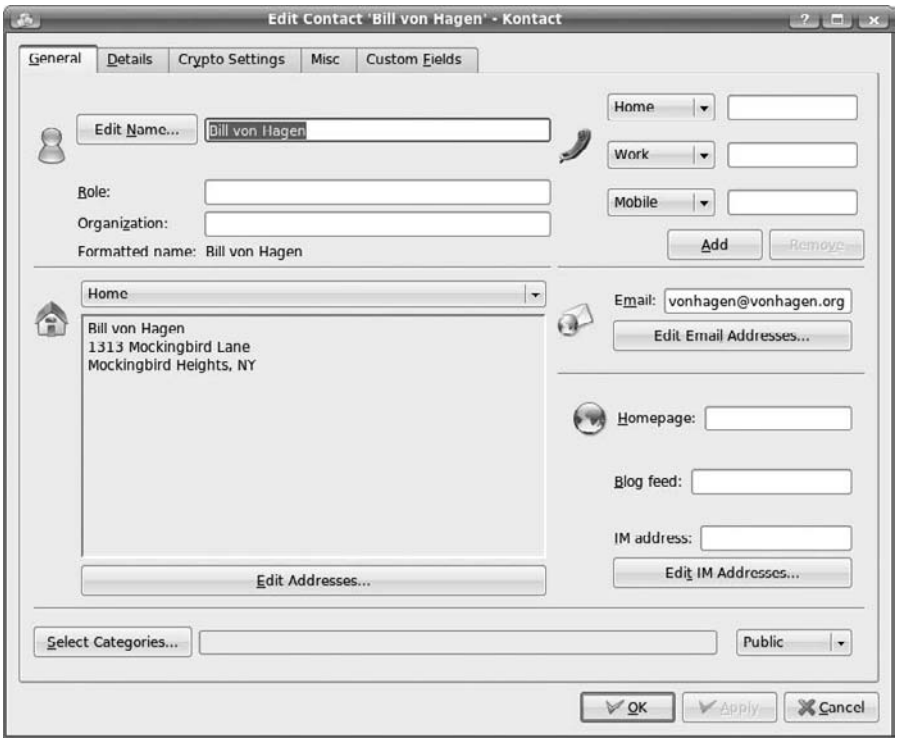


FIGURE 10-43

Creating a new contact in Kontact



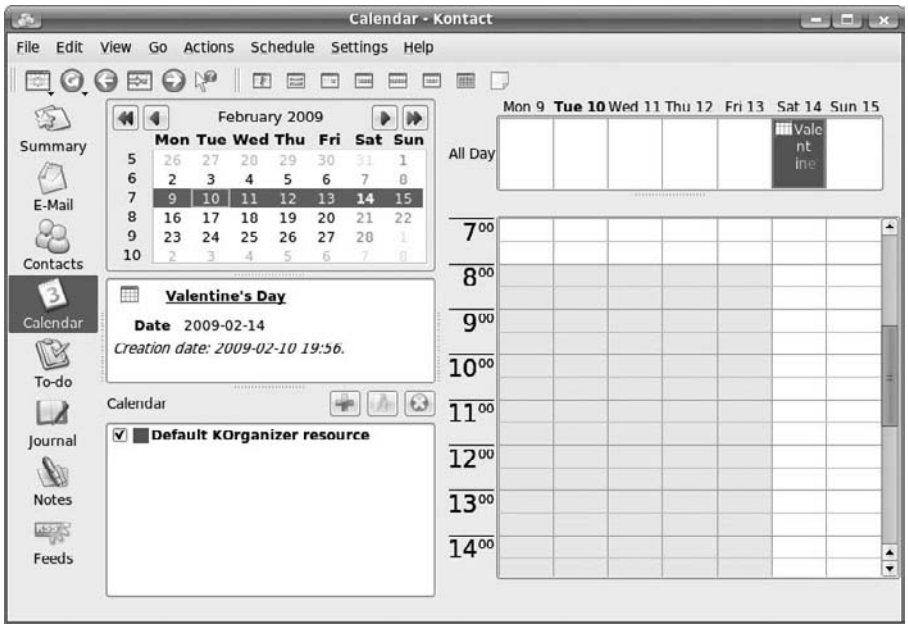
In addition to maintaining local contact information on your system, you can also integrate the Contacts application with networked LDAP servers that provide address book information. For information about identifying these servers, see the section earlier in this chapter entitled "Configuring Contacts."

Using and Managing Your Calendar

Managing one's calendar and To-Do items is an important part of a personal information management application because remembering where you have to be and what you have to do is a fairly critical aspect of getting organized in the first place. As mentioned earlier, managing your calendar, To-Do list, and even a journal is supported in Kontact through an integrated application known as KOrganizer, which is simply known as the Calendar application in Kontact. Kontact's Calendar application lets you manage all of these things locally and, like the Contacts application, also supports integration with centralized calendar resources. For information abut configuring the Calendar application, see the section earlier in this chapter entitled "Configuring Your Calendar".

FIGURE 10-44

A Calendar dialog in Kontact



To access your calendar, click the Calendar icon in Kontact's left pane. A Calendar dialog like the one shown in Figure 10-44 displays.

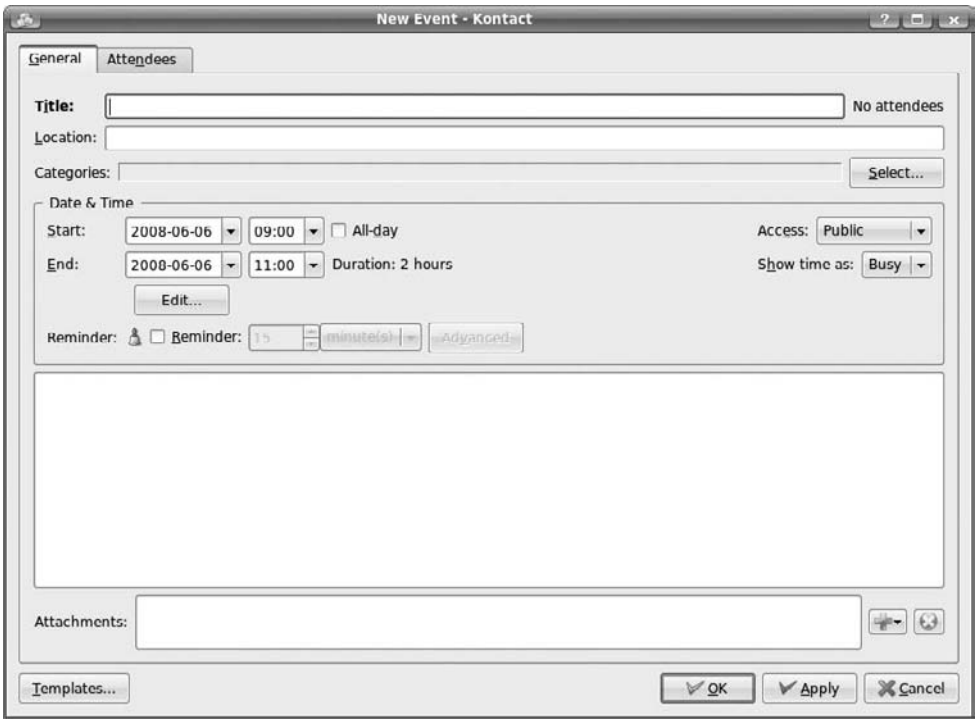
This dialog displays a calendar that you can use to select dates for calendar entries, a summary of selected events or entries, a weekly summary of your current calendar, and a focus view of calendar events for the currently selected day.

To add a new calendar entry on the current day, double-click on the time at which you want to create the entry. To add a calendar event for another day, select that day from the calendar (or the weekly summary, if that date falls in the current week) and double-click on the time at which you want to create the entry. To schedule calendar events further in the future, you can use the single right arrow at the top of the calendar to move forward one month or use the double right arrow to move forward one year. If you overshoot the month you're looking for or want to return to the present date, you can use the left arrow to move back one month or the double left arrow to move back one year.

After selecting the time at which you want to schedule a calendar entry, a dialog like the one shown in Figure 10-45 displays.

FIGURE 10-45

Scheduling a Calendar event in Kontact



You must enter a title for the event in order for it to show up on your calendar in some meaningful fashion. All of the other fields on this dialog are optional, although you may need to configure the end time of the event if it is other than the default of 2 hours. To set a reminder for this event, click the Reminder checkbox and use the number and units (minutes, hours, days) to configure how long before the event you will be reminded.

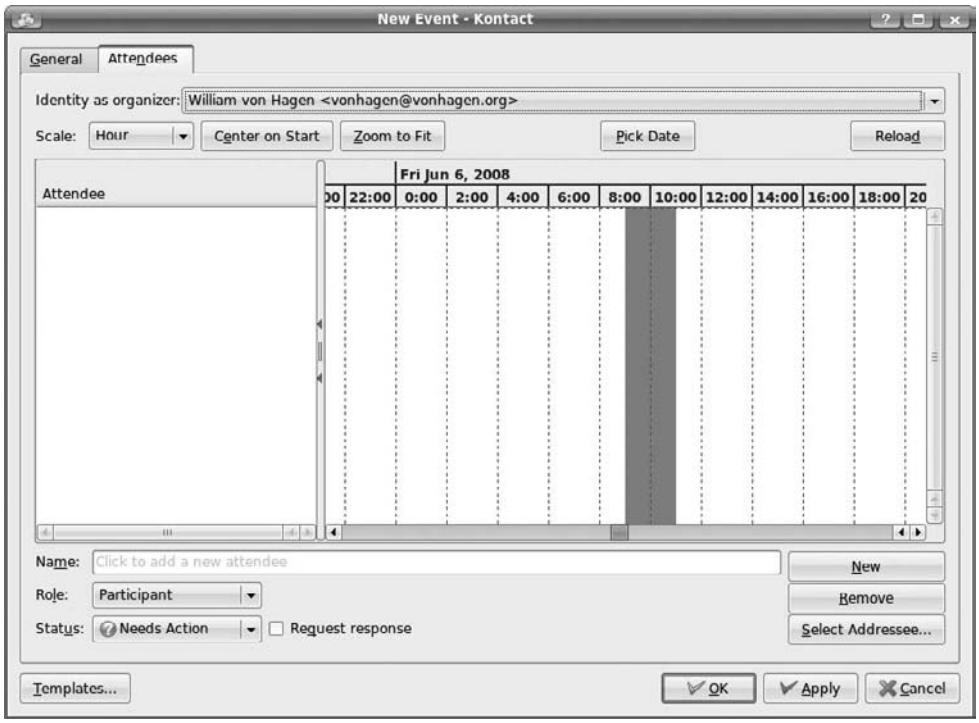
The other fields in this dialog are useful as notes to yourself, and especially if you want to invite other people to this event. To do the latter, click the Attendees tab to display the dialog shown in Figure 10-46.

To add other attendees, click the Name text entry area and begin entering a name or e-mail address. Kontact will attempt to auto-complete the name or address that you enter from the entries in your Contacts list.

Once you've finished entering a Calendar item, click Apply to save your changes and continue editing the current event, or click OK to close the configuration dialog for this calendar event. After inviting multiple attendees to a calendar event, you may want to click Apply and return to the General tab for this calendar event if you want to add e-mail attachments, add or double-check location information, or add a description of the event. I generally do this because I typically find that I've missed something at this point, but that could just be me.

FIGURE 10-46

Inviting other people to a Calendar event in Kontact



TIP **If you're a Google calendar user, see the excellent Linux.com article on integrating the Google Calendar with the Kontact Calendar at** `www.linux.com/feature/122054`**. This article also discusses integrating Gmail with Kontact, which is touched upon in this chapter and is discussed in more detail in the section of Chapter 9 entitled "Accessing Free and Web-Based E-mail."**

Summary

The Kontact Personal Information management suite provides a number of different applications in a single, coherent framework. The fact that all of these applications were formerly stand-alone applications (and can still be executed in that fashion if you create menu items for them or run them from the command line) is a testament to their longevity and stability. Unifying them in the single Kontact framework is the icing on the cake.

At a minimum, almost everyone needs e-mail nowadays, and Kontact provides a powerful and ultimately configurable e-mail application. As you use Kontact, you will probably also find that other Kontact components, such as the Calendar and To-Do lists, are extremely useful. I don't know anyone who couldn't be better organized, and Kontact can help. While I'm a long-term fan of the Evolution PIM that is provided on standard Ubuntu systems, Kontact is an excellent alternative. After seeing Kontact, some of my standard

Ubuntu fans have taken to installing and using Kontact on their GNOME systems, and now swear by Kontact. Part of the beauty of Linux and open source software in general is the ability to pick and choose the tools that work in a way that best suits you, typically at the same high level of quality—and all at the same low price of free.

The next chapter discusses how to surf the Web using Firefox, one of the most popular Web browsers available today. Although Konqueror is the default Web browser provided on Kubuntu systems, installing Firefox is well worth considering due to its power, popularity, and the tremendous number of plug-ins, themes, and other add-ons that are available.

Chapter 11

Surfing the Web with Firefox

As with other open source software, many Web browser options are available to Ubuntu users. Desktop systems such as GNOME and KDE both have multipurpose object browsers (Nautilus and Konqueror, respectively) that can also browse Web pages and are fine for that purpose if you want to learn only one tool and use it for almost everything. However, Ubuntu installs the Mozilla project's Firefox browser as its default Web browser for a good reason — it's the best Web browser available on Linux systems (and on any other system, too, in my opinion). Firefox is the most popular open source browser available today, and it's the only browser in the past decade or so to have actually taken market share away from Microsoft's Internet Explorer. Many of the innovations that Windows users are looking forward to in upcoming releases of Internet Explorer were first introduced in Firefox and have been present for quite a few versions.

This chapter begins with a bit of history and explains what a long, strange trip it has been for Firefox, and then discusses how to start, configure, use, and customize Firefox.

IN THIS CHAPTER

- Reviewing the history of Firefox
- Using Firefox
- Configuring Firefox
- Working with Bookmarks
- Extending Firefox

A Quick History of Firefox

Firefox has what is perhaps the best possible pedigree in Web browsing history — it is the latest descendant of the original Mosaic Web browser, the original graphical Web browser for UNIX systems. Mosaic was created in 1992 at the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign.

TIP **For a quick history of the World Wide Web, see the section of Bonus Chapter 8 titled "World Wide Web 101."**

As the Internet boom unfolded around the Web and its killer application, the graphical Web browser, much of the team that developed Mosaic, including

the leader of the team, Marc Andreessen, left the NCSA to found Mosaic Communications Corporation. Because of issues with NCSA over the Mosaic name, Mosaic Communications Corporation quickly became Netscape Communications Corporation, and the company's fledgling browser product became known as Netscape Navigator. Early versions of Netscape Navigator were available by mid-1994, with the first official release occurring near the end of that year.

NOTE **In one of the incestuous twists of fate that make hi-tech history interesting, a company named Spyglass, Inc. licensed the Mosaic technology and trademarks from the University of Illinois and used it to develop its own Web browser, Spyglass Mosaic, for personal computer systems running Microsoft Windows. Spyglass then eventually licensed the source code for Spyglass Mosaic to Microsoft, where it became the basis for Microsoft's Internet Explorer.**

Netscape Navigator was merely a Web browser, and the growing legions of Internet devotees clamored for other Net-related applications. Demand for other Net-related applications, such as mail clients and Web page composition tools, led Netscape to make Netscape Navigator a part of their new Netscape Communicator Suite.

Internet Explorer quickly began to take over the browser market from Netscape — after all, it came free straight from the mother ship on every Microsoft Windows desktop system, and its accompanying Web server product, Internet Information Server, which was included free with every Microsoft Windows server system. By 1998, Netscape Communications Corporation was essentially dead as a standalone entity, although it was acquired by America Online (AOL) for a surprising $4 billion.

Netscape's last gasp as a standalone entity was to found the open source Mozilla project, donate the Netscape Communicator source code, and create an accompanying Mozilla Organization. Both were chartered with stewarding the creation of the next-generation Internet suite for Netscape. Mozilla was the original code name for the Netscape browser throughout its creation, and an accompanying red dragon was a popular decoration at Netscape. The Mozilla project, often flamed for disorganization, eventually stabilized and modularized its code base, and continues to form the basis of a Netscape browser that is available at `www.netscape.com` and is still under development today.

Meanwhile, the Mozilla project began to create its own Web browser product. This browser was originally named Phoenix, and then Firebird, both of which had naming conflicts with existing applications, so Mozilla eventually settled on the name Firefox.

Throughout its history, Firefox has introduced features that we all now take for granted, such as an integrated pop-up blocker, a customizable toolbar, a sidebar, integrated searches on the toolbar, graphical cookie management, a download manager, customizable themes, the ability to easily add external browser extensions, and tabbed browsing.

NOTE **Although Firefox and its sibling e-mail client, Thunderbird, are the primary products of the Mozilla Foundation, other related applications are still under development and available. The most interesting of these is the SeaMonkey project, a suite of Internet applications based on the original Mozilla Application Suite, the descendant of the Netscape Communicator. The SeaMonkey project includes a Web browser, an e-mail and newsgroup client, an IRC chat client, and a tool for composing HTML Web pages. SeaMonkey is available at `www.seamonkey-project.org`.**

Starting Firefox

Firefox is installed by default on standard GNOME-based Ubuntu desktop distributions, and is one of the stellar examples of open source software quality and design.

NOTE **If you are running Kubuntu, you will have to install Firefox manually. You can do this by installing the firefox *package*, which requires the *firefox-3.0* package and will thus ensure that you get future updates. If you are running Kubuntu and want to get all of the Firefox customizations discussed in this chapter, such as the Get Ubuntu Extensions customization that is available from the Add-ons dialog and other time-savers, you will also have to install the *ubufox* package. Installing the ubufox package will also install a fair amount of GNOME dependencies, but that's worth it in my opinion, because disk space is cheap and my time (and probably your time) is not. As always, your mileage may vary. For information about installing software on Kubuntu systems, see the section of Chapter 23 entitled "Using the Adept Manager to Add and Remove Software."**

You can start Firefox manually in a variety of different ways on a standard Ubuntu system:

- By clicking the Firefox Web Browser icon in the GNOME toolbar at the top of the screen on Ubuntu desktop systems, shown in Figure 11-1. This is the easiest, most common, and most intuitive way to start the browser.
- By selecting Firefox Web browser from the Applications ➪ Internet menu.
- By typing the `firefox` command in a GNOME Terminal or xterm window and pressing Return.

Ubuntu's GNOME desktop is also preconfigured to start Firefox for you if you click on a Uniform Resource Locator (URL) in an e-mail reader or another Internet-aware application.

On Kubuntu systems, you can start Firefox from the Internet menu. On my Kubuntu systems, I typically add an icon for Firefox to the kicker's Quick Launcher by right-clicking on the Quick Launcher and selecting the Internet Firefox Web browser menu command.

NOTE **Installing Firefox on a Kubuntu system does not change the default association of URLs with Konqueror — in other words, clicking on a URL in a Kubuntu application will still start Konqueror.**

The first time you start Firefox on an Ubuntu system, it will automatically start its Add-ons manager and install two language packs — one for Firefox itself, and a second for Xulrunner, which is a run time environment for Mozilla's XUL (XML User Interface Language) Web applications. You can close this dialog for now — see the section entitled "Enhancing Firefox" later in this chapter for more information about installing other extensions, themes, and plug-ins.

FIGURE 11-1

The Firefox icon on your GNOME desktop



After you start Firefox for the first time, a screen such as the one in Figure 11-2 displays, showing the default home page, which displays a "welcome" message and provides general information about the Ubuntu release that you're running. If you've started Firefox for the first time by clicking on a URL in another application, the Web page corresponding to that URL displays.

FIGURE 11-2

Starting Firefox for the first time



Most people configure Firefox (or any Web browser, for that matter) so that it starts up on a Web page of their choosing, referred to as their home page. I'll explain how to configure that later in this chapter — for now, let's explore the Firefox user interface.

The Firefox User Interface

Although most Web browsers look and work approximately the same nowadays, this section provides an overview of the different portions of a Firefox window, just in case you're lucky enough to be running Ubuntu on your first computer. If you're already comfortable with Web browsers and graphical applications in general, feel free to skip this section entirely or skip ahead to the section titled "Special and Not-So-Special Firefox Features" to get an overview of some of the special capabilities available in Firefox that may not have been available in Web browsers that you've used previously.

Standard Parts of a Firefox Window

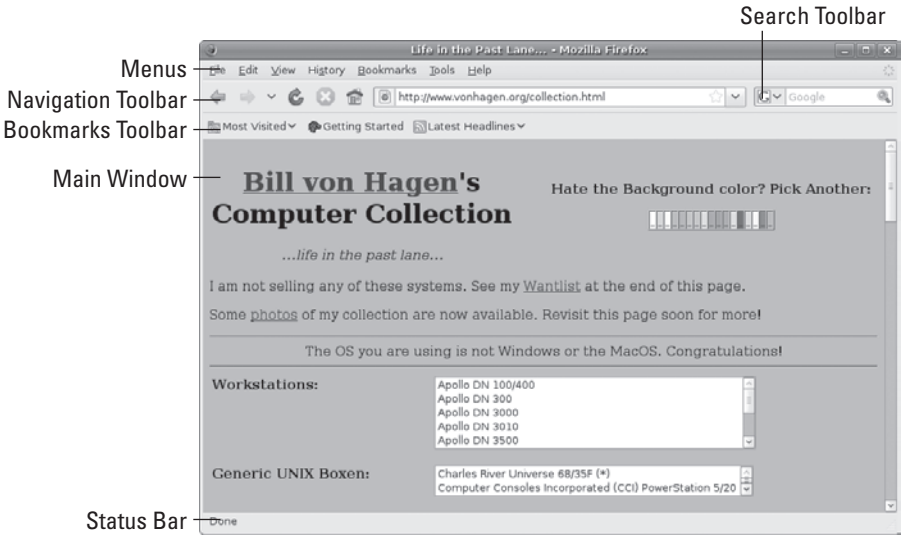
Figure 11-3 shows a sample Firefox screen, showing the different portions of a standard Firefox window.

These different areas provide the following functionality:

- **Menus:** Displays a drop-down menu of related commands when you position the cursor over any of the items in this area and click on them.
- **Navigation Toolbar:** Displays icons that control the Web page displayed in the main Firefox window. By default, the Navigation toolbar also provides quick access to the Google search engine.

FIGURE 11-3

Different portions of a Firefox window



- **Bookmarks Toolbar:** Provides quick access to selected bookmarks that are more frequently visited than the standard bookmarks saved on the Bookmarks menu. In Firefox 3, the Bookmarks toolbar also includes a *Smart Bookmarks* item that provides quick access to recently visited sites (much like the History menu), sites that you visit often, sites that you've recently bookmarked, sites that you've tagged, and so on.
- **Tabs:** Represent different Web pages that you are currently visiting. By default, this portion of the screen displays only if you are currently visiting more than one site, although this behavior is configurable.
- **Main Window:** Displays the contents of the Web page that you are currently visiting.
- **Search Toolbar:** Displays commands that enable you to search for specific text in the Web page that you are currently visiting. This portion of the screen displays only if you have first executed a Search command.
- **Status Bar:** Provides summary information about Firefox's retrieval of the Web page that you are currently visiting. Icons at the right end of this page display status information about that page if indexing is enabled, an RSS feed is available on that page, the page is retrieved through a secure connection, and so on.

As I explain later in this chapter, you can configure many of these items — what they contain (Navigation toolbar, Bookmarks toolbar), how they look (by using custom themes), and whether they display at all (Navigation toolbar, Bookmarks toolbar, status bar).

Using the Mouse in Firefox

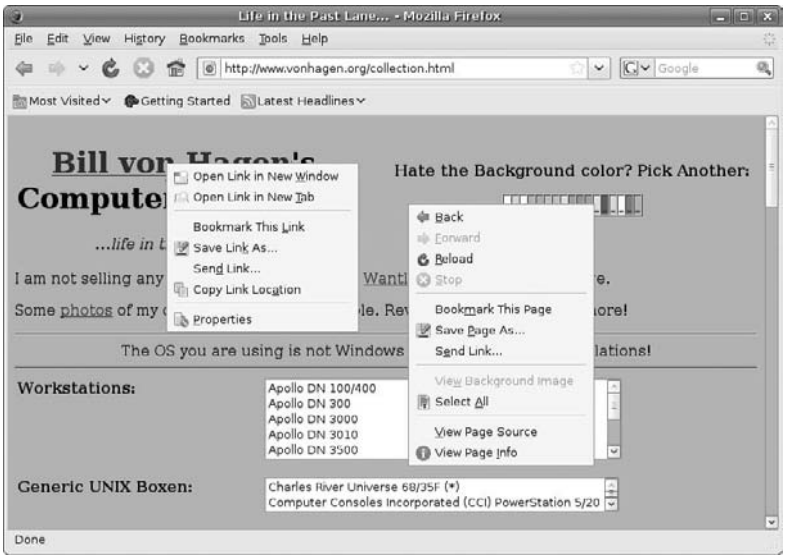
There's nothing different about using the mouse in Firefox than in any other graphical application except for the contents of any menus that pop up in response to a right-click. However, if you're new to graphical applications, let's review. On a standard three-button mouse, clicking (pressing and releasing) a mouse button does the following:

- **Left button:** Selects something. Left-clicking on a menu displays that menu and left-clicking on an item from that menu selects a command or submenu. Left-clicking on a button in the Navigation toolbar executes the command associated with that button or, in the case of the Google toolbar, positions the cursor so that you can specify what you want to search for. Left-clicking on a portion of a Web page performs any action associated with that area — following a hyperlink, displaying a larger graphic, submitting a form, and so on.
- **Middle button:** Depending on where you click, middle-clicking can open a hyperlink in a new tab, paste selected text into a selected text area, and so on.
- **Right button:** Displays different menus depending on what you right-click on. These are known as *context-sensitive menus* for this reason. Figure 11-4 shows two context-sensitive menus from Firefox: The menu at the right displays if you right-click on the background of the current Web page, while the menu at the left displays if you right-click on a hyperlink in the current Web page. (By the way, you can't get both of these context menus to display at the same time as this figure might suggest.)

You can customize context-sensitive menus such as those shown in Figure 11-4 by adding Firefox extensions that have associated menu items, as explained in the section "Adding Firefox Add-ons and Themes," later in this chapter. You can even customize how the mouse behaves to some extent by adding Firefox extensions that tell Firefox to perform specific actions when specific types of mouse movements occur.

FIGURE 11-4

Context-sensitive menus from Firefox



Special and Not-So-Special Firefox Features

You may be familiar with the vanilla Web browsers installed with other operating systems, such as Internet Explorer on Microsoft Windows and earlier releases of Mac OS X. Firefox provides the same capabilities as those browsers and adds a few powerful, exciting new ones. Some of the most interesting features first added by Firefox and standard browser capabilities that aren't active by default in Firefox are as follows:

- **Tabbed browsing** enables you to open multiple Web pages simultaneously without creating separate windows for each. Each Web page that is currently open is represented by a separate tab, which displays above the portion of the screen that shows the current Web page. You can quickly and easily switch between these Web pages by clicking on the appropriate tab. This feature is so useful that it has been incorporated into most other modern browsers — for example, Microsoft added this capability in Internet Explorer 7, and this feature is also available in Mac OS X's Safari browser. To close a selected tab, click the red X icon at the top right of the Tab display area.
- **Sidebars** provide easy navigation through a given list of Web pages. Firefox provides two different sidebars, one displaying your bookmarks and a second that provides a list of the Web pages that you have visited recently (known as a *history* sidebar). These sidebars aren't turned on by default in Firefox, but you can easily activate whichever you prefer by selecting Bookmarks or History from the View ➪ Sidebar menu. Figure 11-5 shows Firefox with the Bookmarks sidebar displayed. To close a sidebar, click the red X icon at the top right of the sidebar.

FIGURE 11-5

The Bookmarks sidebar in Firefox



- **Missing Plugin Detection** automatically identifies sites that provide content that requires that you install a special plug-in in order to view that content. Common examples of this are

multimedia content types such as Adobe Flash content, Apple QuickTime content, and so on. Figure 11-6 shows the banner that displays (below the tabs) when Firefox identifies a missing plug-in. You can click the Install Missing Plugins button in this banner to have Firefox start its Plugin Finder Service, which will attempt to locate and enable you to install plug-ins that can handle the requested type of content. If this banner displays and you don't want to change anything, you can close it by clicking the red X icon at the right of this banner. See the section of this chapter entitled "Adding New Firefox Plug-ins" for more information about this service.

FIGURE 11-6

The missing plug-in banner in Firefox



- **Built-in Popup Blocking** prevents sites from displaying annoying pop-up windows unless you specifically authorize a site to display them (or you disable this feature entirely). You can click the Preferences button in this banner to customize how Firefox handles pop-ups from the current site, or in general. If this banner displays and you don't want to change anything, you can close it by clicking the red X icon at the right of this banner. Customizing how Firefox handles pop-ups is explained later in this chapter.
- **Live Bookmarks** provide a special type of bookmark that is dynamically constructed from an RSS feed provided by a specific site. Creating Live Bookmarks is explained later in this chapter.

Firefox is a powerful, flexible Web browser with all of the features that you'd expect in any modern Web browser, and then some. You could easily devote an entire book to exploring each and every menu and Firefox feature, but my publisher probably prefers that this not be that book. So I'll agree, and just encourage you to explore the various Firefox menus and the hundreds of available themes and extensions (some of which are discussed later in this chapter).

Configuring Firefox

Firefox is extremely configurable, and provides enough knobs to turn that you could probably spend as much time tweaking your configuration and enhancing Firefox as you can in actually using the browser. Luckily, books such as this one and Firefox's own documentation can help make sure that you spend more time browsing than fine-tuning. Firefox provides great online help at `www.mozilla.org/support/firefox/`, which includes a tutorial that will walk you through its most popular configuration tasks.

Almost all of the permanently configurable aspects of Firefox are located in its Preferences dialog, which you display by selecting the Edit ➪ Preferences menu item. Figure 11-7 shows the Firefox Preferences panel in the version of Firefox that was current when this book was written. Yours may look slightly different, but the core contents are the same.

FIGURE 11-7

The Preferences window in Firefox



This section discusses the highlights of Firefox configuration, focusing on configuration options that you'll eventually want to know about, regardless of whether or not you want to customize them immediately.

Setting Your Home Page

Your *home page* is the Web page that displays when you first start Firefox (or any Web browser). Many people use a Web portal site of some sort, such as yahoo.com, msn.com, aol.com, or netscape.com, as their home page; others use news sites that are specific to their job or interests, and others use personally created Web pages on Internet sites that they have created or are affiliated with. Regardless of what you use, setting your home page is one of the key aspects of personalizing your use of Firefox.

To set your home page in Firefox, first display the Preferences dialog. The Main pane should be displayed by default. If it is not, click the Main button at the top of the dialog to display the pane shown previously in Figure 11-7.

The Startup section at the top of this pane enables you to configure the page that displays when Firefox starts. Clicking the aptly named "When Firefox starts" drop-down gives you three general options:

- **Show my home page:** Displays a specific Web page that you identify as your home page
- **Show a blank page:** Displays a blank page for faster startup
- **Show my windows and tabs from last time:** Displays the active page and all of the tabs that were visible the last time that you exited Firefox

I typically use a home page that lets me check for important news (and not-so-important news) before moving on to whatever specific sites that I want to browse. Firefox offers four options for setting your home page:

- Explicitly typing a URL into the Location text area
- Clicking Use Current Page to use the page that is currently displayed in Firefox
- Clicking Use Bookmark, selecting one of your existing bookmarks from the dialog that displays, and clicking OK to close that dialog
- Clicking Restore to Default to cause Firefox to display the default "Welcome to Ubuntu" screen shown in Figure 11-2

Once you've chosen one of these options to select your new home page, click Close to dismiss the Preferences dialog. The home page that you specified will be displayed the next time you start Firefox.

Controlling Pop-ups, JavaScript, and More

Firefox provides a single Preferences dialog for configuration, so in the same place where you selected your default home page, you can also configure Firefox's support for various features of modern Web pages, and how these Web features behave when enabled. You can enable, disable, or customize the behavior of the following from a single dialog:

- **Block Pop-up Windows** determines whether to block all remote sites from displaying pop-up windows, allow selected sites to display pop-ups, or allow all sites to display pop-ups on your system.
- **Load Images Automatically** determines whether to load all images on any Web page, load images only from the site that actually hosts the Web page (to block ads and graphics provided by other sites), and load images from a list of sites even if loading images is prohibited.
- **Enable Java** determines whether to support Java in Firefox so that remote sites can download Java applets that will execute locally on your machine.
- **Enable JavaScript** determines whether to support JavaScript, which enables programmatic content generation and window manipulation, and (if so) determines which specific features to support.

To enable, disable, or customize these content-related features, first display the Preferences dialog. Next, click the Content button at the top of the dialog to display the pane shown in Figure 11-8.

FIGURE 11-8

Enabling and customizing content in Firefox



Once this pane displays, you can select the checkbox beside the feature you're interested in tuning to enable or disable that activity. For options that allow you to create lists of sites to be excluded from a specific setting, click the Exceptions button to the right of that option, and enter the names of the sites that you want to exclude from that setting. Similarly, support for JavaScript is usually enabled by default in Firefox, but the dialog that displays after clicking the Advanced Settings button gives you very fine-tuned control over the sorts of things that JavaScript-enabled Web pages can do on your system.

Once you've finished configuring these content-related features, click Close to dismiss the Preferences dialog. Changes to some settings will not take effect until you restart Firefox, so it's always a good idea to exit from and restart Firefox after changing content settings, just in case.

Configuring Your Privacy Settings

In the Privacy pane of the Preferences dialog, different topics are organized onto separate tabs across the center of the pane. Customizable privacy settings include the following:

- **History** determines the period of time (in days) for which Firefox keeps a record of the URLs for the Web sites that you've visited. This also includes a setting that determines whether Firefox keeps a list of the files that you have downloaded.
- **Cookies** determines whether or not Firefox accepts cookies in general, a list of sites excluded from this policy, and, if so, how long cookies are retained. Cookies are encoded chunks of information remote Web sites can set in your browser to immediately recognize you the next time

you visit. Like the Saved Passwords setting, this configuration item provides a special dialog that enables you to manage the list of cookies that are currently stored on your system.

- **Private Data** determines whether Firefox remembers private data, such as the values that you've entered when searching for things in the search toolbar or when filling out forms on the Web.

NOTE **If you've already been using Firefox for a while, this pane will look very different to you. Older versions of Firefox used a tabbed Privacy pane to control all of your privacy-related settings. Different tabs contained groups of related privacy settings. The current version of Firefox uses a simpler interface, dividing privacy-related settings into two separate panes: the Privacy pane, shown in Figure 11-9, and the Security pane, shown in Figure 11-10.**

To enable, disable, or customize these privacy-related settings, first display the Preferences dialog. Next, click the Privacy button at the top of the dialog to display the pane shown in Figure 11-9.

FIGURE 11-9

Customizing privacy settings in Firefox



Once this dialog displays, click any options that you want to change. For example, if you are using a computer that you share with others, you may want to select the Always clear my private data when I close Firefox option to help ensure that no other user of that system can access your personal information. Many people are also paranoid about cookies, which are used in both good and irritating ways. For example, cookies can store shopping cart information, which is very useful if you are buying something online and have to interrupt your session without making an actual purchase. On the other hand, many

online sites and advertisers use cookies to track what you've looked at on their sites. This can be useful in terms of showing ads related to your buying and browsing decisions, but it also provides personal information that you may not want to share. In my opinion, browsing the Web without using cookies is a big hassle, so I always accept them, but I occasionally use the Show Cookies button on this Preferences pane to browse the list of cookies that are stored on my system and delete those cookies that are for sites that I don't recognize or care about anymore.

After making any changes to the Privacy settings on this dialog, you can either click Close to save your changes and exit the Preferences dialog, or you may want to click the Security icon at the top of the Preferences pane to explore your Firefox security settings, which are certainly related to maintaining your privacy. Clicking the Security icon displays the pane shown in Figure 11-10.

FIGURE 11-10

Customizing security settings in Firefox



The first few entries on this pane enable you to configure whether remote sites can install Firefox add-ons, whether Firefox will complain when you attempt to visit a site that is believed to contain spurious software, and whether Firefox will warn you if it believes that you are being redirected to a site other than the one that you requested. The Passwords section determines whether or not Firefox stores username/password pairs that you've used to log in to different secure sites. This configuration item provides special dialogs for managing any login/password combinations that are currently saved. Clicking the Saved Passwords button displays the Saved Passwords dialog shown in Figure 11-11, which enables you to explore and manually delete any username/password pairs that are currently saved in Firefox.

FIGURE 11-11

Managing passwords saved in Firefox



To remove a saved username/password pair for a specific site, highlight that entry and click Remove at the bottom of the dialog, as my wife would probably like to do with my eBay login information. As with most graphical selections, you can hold down the Ctrl key while clicking on entries to select multiple username/password pairs, or select one item, hold down the Shift key, and click another to select those two items and all of the items between those two username/password pairs. Click Close to dismiss this dialog and return to the Privacy Settings pane when you're done.

Once you've finished configuring your privacy and security settings, click Close to dismiss the Preferences dialog. Changes to some settings will not take effect until you restart Firefox, so it's always a good idea to exit from and restart Firefox after changing privacy settings, just in case.

Working with Bookmarks

In a Web browser, a bookmark is a record of the Uniform Resource Locator (URL) for a Web page that you've visited, and which you want to visit again. When you select a bookmark, Firefox goes directly to the page that it represents — no fuss, no muss.

NOTE **"Bookmarks" is the more common term for what Internet Explorer calls Favorites.**

Bookmarks are great things, just as they are in real books. Without a bookmark, you'd have to start searching a book at the beginning each time you picked it up. The same thing is true of bookmarks in terms of Web searches — once you've found somewhere that you want to return to on the Web, it's important to save a bookmark to that location so that you can easily return there. You certainly don't

want to have to redo a complex Web search process or Web page navigation sequence any more times than you have to — which should be once, if you've bookmarked a page.

Firefox provides two primary locations in which to save bookmarks: a hierarchical collection of bookmarks that you can access by selecting the Bookmarks menu, and the Bookmarks toolbar, which is displayed directly below the Navigation toolbar in the main Firefox window. The Bookmarks toolbar is designed to provide quick access to frequently visited bookmarks, while the Bookmarks menu provides access to everything. You can think of the Bookmarks toolbar as the bestsellers section of a bookstore, while the Bookmarks menu is the entire store.

Over time, most people tend to accumulate huge collections of bookmarks, which makes managing them — and being able to find the one you want when you want it — a very critical task for continued happiness through Web surfing. Firefox includes an excellent Organize Bookmarks dialog that simplifies both of these tasks. The next few sections explain how to create bookmarks in various ways, how to organize them using the Organize Bookmarks dialog, how to back up your bookmarks and even access them from other Web browsers, and how to use an exciting feature introduced by Firefox, known as Live Bookmarks, to automatically construct dynamic menus that provide an overview of selected Web sites.

Creating Bookmarks in Firefox

Firefox provides three basic ways of creating a bookmark to the page that you're currently viewing:

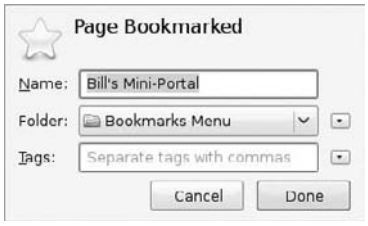
- Click the star to the right of the Web site's URL in the Firefox navigation bar. After doing so, you can click the star again (now yellow) to specify the location to which you want to save the bookmark, add tags, and so on.
- Press Ctrl+D on your keyboard (hold down the Ctrl key and press the letter *D*).
- Right-click on the background of the current page and select the Add Bookmark command from the pop-up menu.
- Select the Add Bookmark command from the Bookmarks menu.

You can also create a bookmark for a link on the page that you're viewing by right-clicking on the link and selecting the Bookmark This Link command from the pop-up menu.

Regardless of the method you use to create a bookmark, creating one displays the dialog shown in Figure 11-12. The top portion of the dialog shows the name that the bookmark will have. By default, this is the text under the link or the name of the current page, but you can change it to be something more memorable if you'd like.

FIGURE 11-12

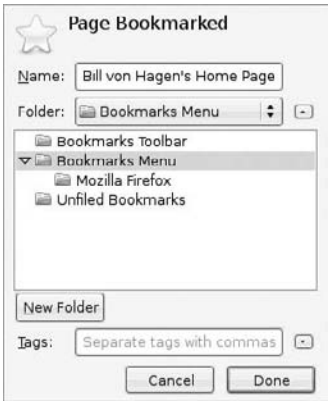
The Add Bookmark dialog in Firefox



The Folder field of this dialog identifies where the bookmark will be saved. You can click on this entry to see a list of the folders that already exist in your bookmarks, or you can click the down arrow to the right of this dialog to display a browseable dialog that enables you to select a folder within those folders, as shown in Figure 11-13, or click on the New Folder button to create a new folder in which to store your new bookmark.

FIGURE 11-13

Navigating folders when saving a bookmark



The dialog shown in Figure 11-12 also provides a field that enables you to add tags to your bookmarks. Tags are a type of metadata, which means "information about information." In this case, tags enable you to associate keywords with specific bookmarks so that later you can search for related items. To add tags to a bookmarked entry, enter them as comma-separated keywords in the tags field on the bookmarks dialog, as shown in Figure 11-14.

FIGURE 11-14

Adding tags when saving a bookmark



When adding subsequent bookmarks, you can click the down arrow to the right of the Tags field to see the tags that you've already used in other bookmarks. This can help you use a consistent set of tags across your bookmarks, simplifying future searches.

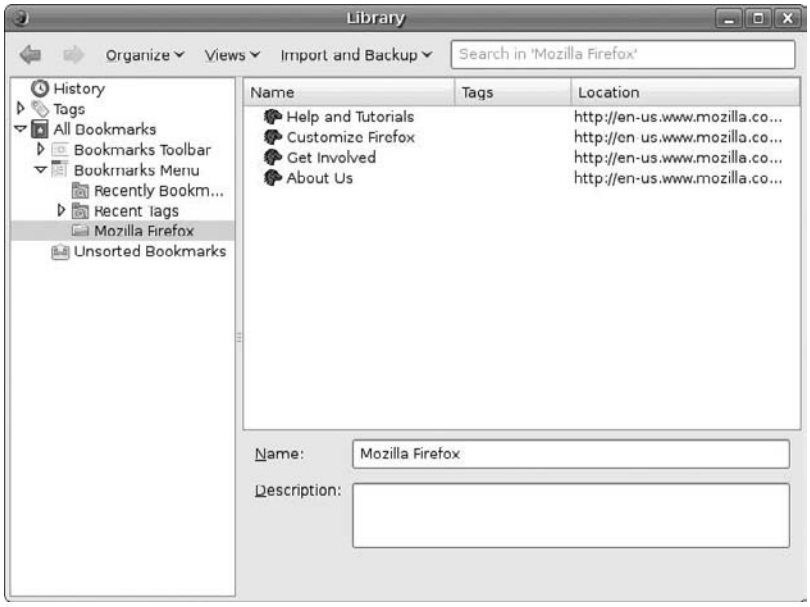
Once you've made any modification that you want to the name of the bookmark and selected the location where you want to store it, click Done to save the bookmark in that location. The dialog closes, and you can now easily go to that page again by selecting the bookmark.

Managing Your Bookmarks

Over time, most people tend to accumulate huge collections of bookmarks, which makes it harder and harder to find the right bookmark — which somewhat defeats the idea of having bookmarks in the first place. To solve this problem, Firefox includes an excellent bookmarks organizer that makes it easy to organize and search your valuable collection of bookmarks, known as your Library. The Library provides a central tool for managing your browsing history and bookmarks, generally referred to as Places. To display the Library, select Organize Bookmarks from the Bookmarks menu, and the dialog shown in Figure 11-15 displays.

FIGURE 11-15

The Library dialog in Firefox



The panel on the left provides a hierarchical overview of how your bookmarks are currently organized, while the panel on the right provides a detailed view of the bookmarks in the folder that is currently selected. You can move bookmarks around by dragging and dropping them (left-clicking on the bookmark and dragging it to the new location) within either panel or even between the two. You can create

new folders, new bookmarks, or new separators (the lines that appear in your bookmarks menu) by selecting the appropriate location in either panel and doing either of the following:

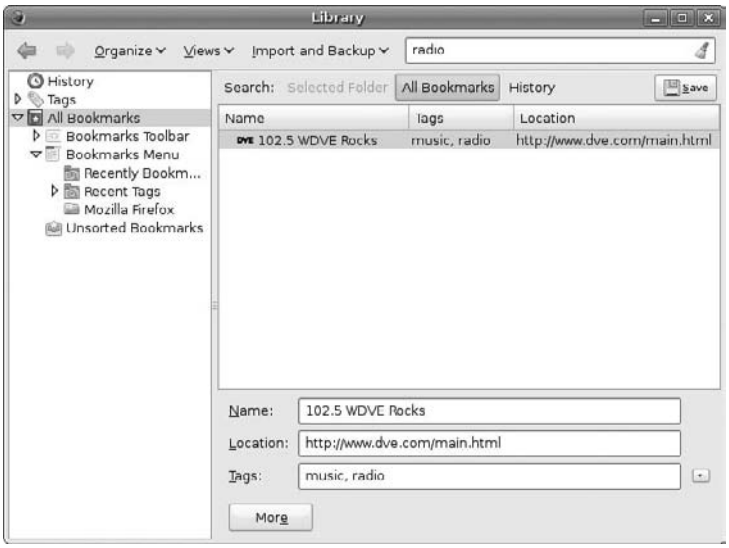
- Select the appropriate icon from the top of the Library dialog.
- Right-click and select the appropriate command from the pop-up menu.

Any changes that you make are immediately reflected on the Firefox Bookmarks menu and Bookmark toolbar.

The Library dialog also provides a convenient way to search your existing collection of bookmarks. You can search your bookmarks for text that appears in their name or URL, but you can also search your bookmarks based on any tags that you've applied to them. Figure 11-16 shows the result of a search for the tag "radio," which I generally apply to the bookmarks for local and Internet radio stations that I'm interested in.

FIGURE 11-16

Searching for bookmarks by tag



Once you're finished reorganizing or exploring your bookmarks, click the Close box in the upper-right corner of the dialog to close the Library dialog.

Regardless of how you choose to organize your bookmarks, one traditional rule is to avoid having one huge list of bookmarks. You probably have specific interests, and can easily store your bookmarks in folders that represent those interests. This makes it easy to find a given bookmark, and also helps prevent you from accidentally bookmarking the same page eleven times because you couldn't find a specific bookmark when you were looking for it. Regardless of how well you organize your bookmarks, the Library dialog's search capability makes it easy for you to find specific bookmarks, but I don't feel that being able to search for something is an excuse for poor, or no, organization.

Backing Up Your Bookmarks

Because your evolving collection of bookmarks is one of the most precious things that you build up as you surf the Net, making sure that you don't accidentally lose them is critical. The Library dialog's Import and Backup drop-down menu provides both Export and Backup commands that enable you to save a copy of your current bookmarks in a separate file. This is especially important for the latest versions of Firefox, which store bookmarks and history in SQLite database files using an interface known as Places. Places is convenient for searching, but removes some of the ease of use of backing up bookmarks in previous versions of Firefox, where you could simply copy a flat HTML file called `bookmarks.html` from your personal preferences directory in the directory `~/.mozilla/firefox` in your home directory to back up your bookmarks. This file still exists, but is only kept in sync with your SQLite bookmark files in certain cases, and there's no guarantee that this will continue in the future. Use the Library dialog's Export command to save an HTML copy of your bookmarks, or use the Backup command to save a copy of your bookmarks in JSON (JavaScript Object Notation) format.

When creating JSON backups, you can easily restore them using the Choose File command from the Restore menu on the Library's Import and Backup drop-down menu. This command enables you to restore a JSON backup file that you've saved manually. The Restore menu also lists dates for which internal Firefox backups are available.

I tend to save my bookmarks in HTML every week or two, giving them names containing the current date, like `bookmarks-27-Oct-06.html`, so that I can tell one from the other. This makes it easy to find a specific bookmark in one of these files if you've accidentally deleted the bookmark but have some idea of the last time that you saw it. JSON is a fine format, but it isn't supported in other browsers or in older versions of Firefox, and I also like to put a copy of recent bookmarks on my Web site to help me find my favorite bookmarked locations even when I'm not in front of my primary machine.

Creating Live Bookmarks

Live Bookmarks are one of the latest — and most useful — enhancements to Web surfing that have been introduced by the Firefox team. Live Bookmarks exploit a popular Web technology called RSS, which stands for Really Simple Syndication, or Rich Site Summary, depending on whom you ask. Many Web sites today use RSS to publish a summary of recent changes to their site (known as an *RSS feed*), which other sites can read and display to provide links to those changes. A live bookmark to a site is essentially a special type of bookmark that points to the RSS feed on that site, made more usable by the fact that Live Bookmarks display the headings in an RSS feed as menu items.

Sites that provide RSS feeds that can be used as Live Bookmarks in Firefox display a Live Bookmark icon at the far right of their URL in the Firefox Navigation toolbar, as shown in Figure 11-17.

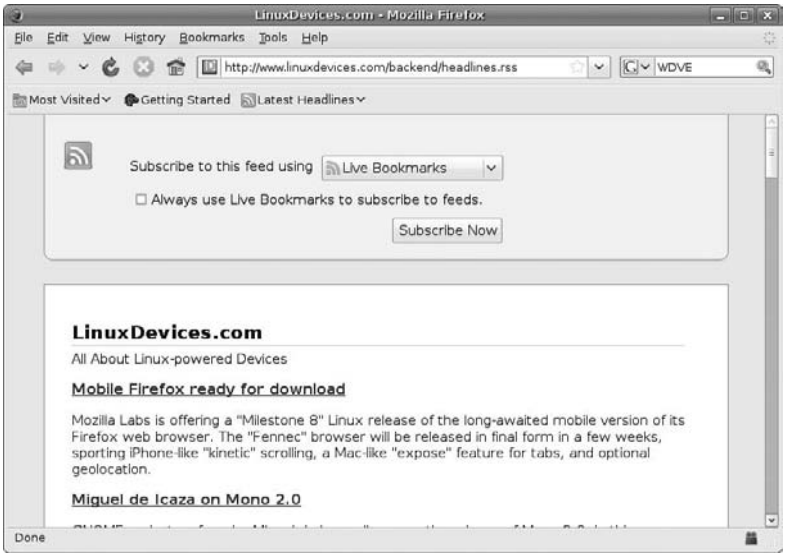
FIGURE 11-17

The Live Bookmark icon for a URL in Firefox



FIGURE 11-18

The Live Bookmark subscription screen

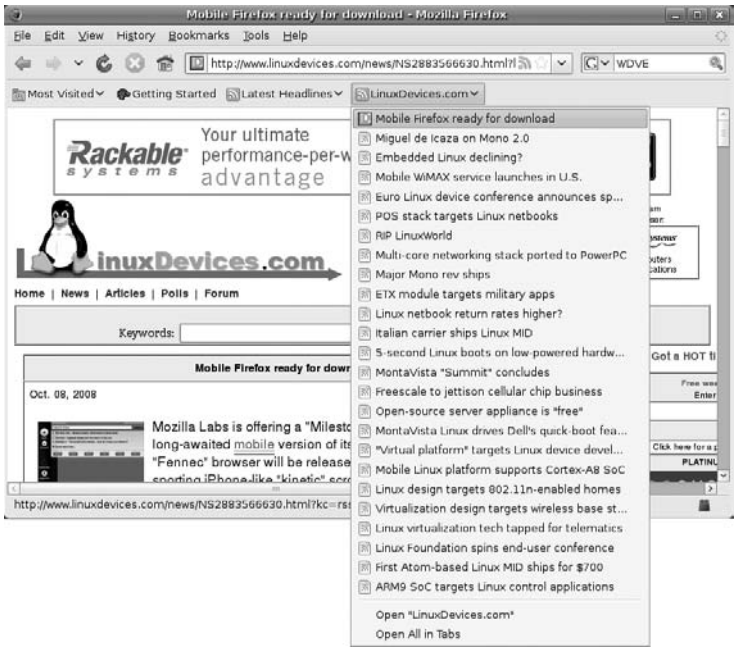


To create a live bookmark, click the Live Bookmark icon in the URL for that site. The screen shown in Figure 11-18 displays.

Click Subscribe Now. The standard Add Bookmark dialog displays, enabling you to specify where you want to create the live bookmark. The Add Bookmark dialog closes, and the live bookmark menu item displays. After a few seconds (the time it takes for Firefox to read the RSS feed and create the live bookmark), you can access the live bookmark menu like any other menu in Firefox, as shown in Figure 11-19.

FIGURE 11-19

A live bookmark menu in Firefox



TIP **Your Bookmarks toolbar is a group of frequently used bookmarks that Firefox displays above the Web page that you are a viewing. This is the most common location for creating Live Bookmarks because they represent an up-to-date summary of a site that you're presumably very interested in.**

The live bookmark shown in Figure 11-19 is for the site that is currently shown in the main browser window in the same figure. If you look closely, you'll notice that the first item on the menu is the first heading on the page in the main browser window. As the site is updated, my live bookmark menu will automatically update to reflect the changes to the site. What could be easier? Perhaps Firefox could also read each of those topics and tell me which ones I should actually care about — maybe in the next release?

Enhancing Firefox

This chapter has stressed the power and usability of Firefox as it comes "out of the box." However, Firefox provides a final set of capabilities that make it even more interesting: You can easily add new Firefox capabilities (known as *add-ons* or *extensions*); new ways of displaying icons, windows, menus, and so on (known as *themes*); and applets (known as *plug-ins*) that run within the Firefox framework and add support for specific types of multimedia files. This section explains how to add each of these items to further customize your Web surfing. In many cases, Firefox extensions provide extra capabilities that are worthy of being separate applications — but which Firefox smoothly integrates into your existing Web surfing experience. Similarly, everyone has their own idea of meaningful icons and appealing graphical design. Firefox themes — much like the skins that you may be familiar with from popular audio applications such as WinAmp or XMMS — enable you to customize the browser's appearance so that it is more attractive to you and quite possibly more usable for you.

NOTE **Old versions of Firefox used separate dialogs for retrieving, installing, and managing extensions and themes, and didn't provide a centralized mechanism for retrieving and installing plug-ins. More recent versions of Firefox provide a single menu command and dialog for performing all of these related tasks.**

Registering for Add-ons

With versions of Firefox prior to 3.0, anyone could download and install any add-on, extension, or theme from their respective Web sites. With Firefox 3.0, some extensions and themes are available only if you've registered with the Mozilla Add-ons site (`https://addons.mozilla.org/en-US/firefox/`). On the one hand, this is a pain and isn't going to help speed the adoption of Firefox; on the other hand, it's useful for the Mozilla Foundation and extension/theme developers to know how many people are using Firefox, and which extensions and themes they're actually using.

You may not want to register until you run across an extension or theme that requires you to do so, but that didn't take me long. You can register with the Add-ons site by clicking the Register link in the upper-right corner of any page, or by going directly to `https://addons.mozilla.org/en-US/firefox/users/register`. After entering your e-mail address, a password, and a minimal amount of personal information, you'll be sent a confirmation message via e-mail. When you receive this, click the link in the message to confirm your subscription, and you're good to go. After registration, Firefox will also display a banner offering to save your login/password information. Even if you're paranoid about saving logins/passwords in browsers in general, saving this one in Firefox seems pretty benign.

After registering and confirming your account, you can log in on the Add-ons Web site at any time by clicking the Log In link in the upper-right corner of any page on that site. This will help make sure that you can download and install any extensions or themes that are available for your version of Firefox.

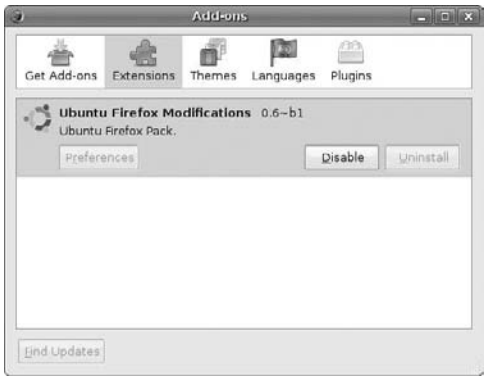
Viewing Installed Firefox Add-ons and Themes

The single dialog used in Firefox for themes, extensions, and add-ons makes it easy to see which of these are currently installed in Firefox, which of those add-ons are currently enabled, and which theme is currently being used. While multiple extensions can be enabled at the same time, only one theme can be used at a time.

To see the add-ons that are currently installed in the version of Firefox that you're using, select the Add-ons command from the Tools menu. This displays the Add-ons dialog. Click the Extensions icon to display the pane shown in Figure 11-20, which shows a list of all of the currently installed extensions. This list will be almost empty initially, but we're going to fix that in the next section.

FIGURE 11-20

The Add-ons dialog in Firefox



NOTE **If you have recently installed new extensions or themes, you will see a yellow banner identifying the number of new add-ons that have been installed. You can close this banner by clicking the button in the right of the banner.**

If an extension is active, you will see a Disable button that enables you to keep the extension installed but deactivate it. If any aspect of the extension is configurable, you will see a Preferences button that enables you to customize the behavior of that extension. If a disabled extension is compatible with your version of Firefox but is disabled, you will see an Enable button.

To see the themes that are currently installed in the version of Firefox that you're using, select the Add-ons command from the Tools menu if you are not already viewing the Add-ons dialog. Click the Themes icon to display the pane shown in Figure 11-21, which displays a list of all of the currently installed themes. As with the list of default Extensions, this list will be almost empty initially, but we'll also fix that in the next section.

FIGURE 11-21

The Themes dialog in Firefox



If you have multiple themes installed, the active theme displays a disabled Use Theme button. All other themes display a Use Theme button, which you can click to enable them. As you click on any theme in the Theme List portion of this dialog, a preview of the icons associated with that theme will display in the right portion of this dialog.

Adding Firefox Add-ons and Themes

Adding new add-ons and themes to Firefox is easy because Firefox will automatically download and install them for you — but figuring out which ones you like best can be hard because there are so many to choose from. But that's what the hours between midnight and 6 a.m. are for anyway, right?

To add new extensions or themes, first select the Add-ons command from the Tools menu. This displays the Add-ons dialog, which initially retrieves and displays a list of recommended add-ons that are not currently installed on your system, as shown in Figure 11-22.

FIGURE 11-22

Recommended add-ons



Once this dialog displays, you can click the Add to Firefox button to install a selected add-on immediately, or you can use either of the links at the top right of the dialog to browse for available extensions (Browse All Add-ons) or browse a list of add-ons that are recommended by the Ubuntu folks (Get Ubuntu Extensions). First, let's explore the huge set of add-ons that are available from Mozilla's Firefox Add-ons site. Click the Browse All Add-ons link to open a new Firefox tab that connects to Firefox headquarters at the Mozilla project and displays a page that shows all of the extensions that are currently available there, as shown in Figure 11-23.

The Categories menu at the left of the Firefox Add-ons page lists a number of topics into which available add-ons are organized, making it easier to find add-ons that are related to some specific task that you want to accomplish in Firefox. These categories can contain both extensions and themes, although there is a separate Themes category that I'll discuss in a moment.

FIGURE 11-23

Available extensions from Mozilla



The search dialogs enable you to search for one or more words in the titles and descriptions of available add-ons. You can restrict your search to a specific category by selecting that category from the drop-down category menu in the search bar. For example, Figure 11-24 shows a search for add-ons containing the word "Theme" in the "Themes" category. (I searched for the word "Theme" because this category also contains appearance-related extensions.)

FIGURE 11-24

Searching for a keyword in a category



You can scroll through this list to see the available add-ons that match your search term. A deactivated button at the right displays if a given add-on is incompatible with your version of Firefox. A green Add to Firefox button displays to the right of any add-on that you can install in your version of Firefox. Figure 11-25 shows themes with each of these buttons (the one at the bottom is the green one).

FIGURE 11-25

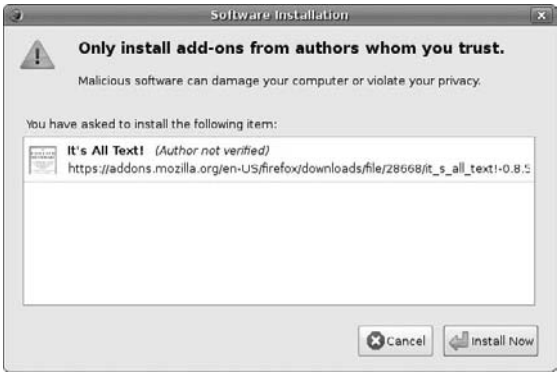
Installable and uninstallable themes



Once you've located an add-on that you want to install, click the Add to Firefox button to automatically download and install the selected add-on. After Firefox downloads the add-on, a dialog like the one shown in Figure 11-26 displays.

FIGURE 11-26

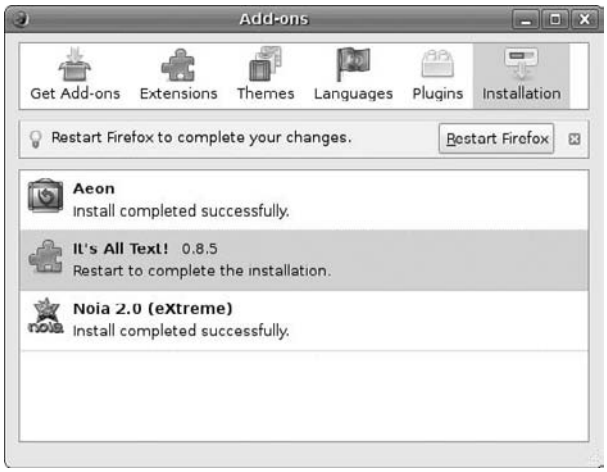
Installing an add-on



Click the Install Now button to install the requested add-on. Once you've downloaded and installed an add-on, an Installation button displays in the Add-ons dialog's toolbar and the associated pane displays to list the new add-on(s) that you've installed, reflecting the fact that the new extension is installed, but will not be active until after you restart Firefox, as shown in Figure 11-27.

FIGURE 11-27

Newly installed add-ons



At this point, you can continue to browse for other themes and extensions at the Mozilla site and install any others that seem interesting. If you're installing more than one extension or theme, you can defer restarting Firefox until you've downloaded and installed all of them.

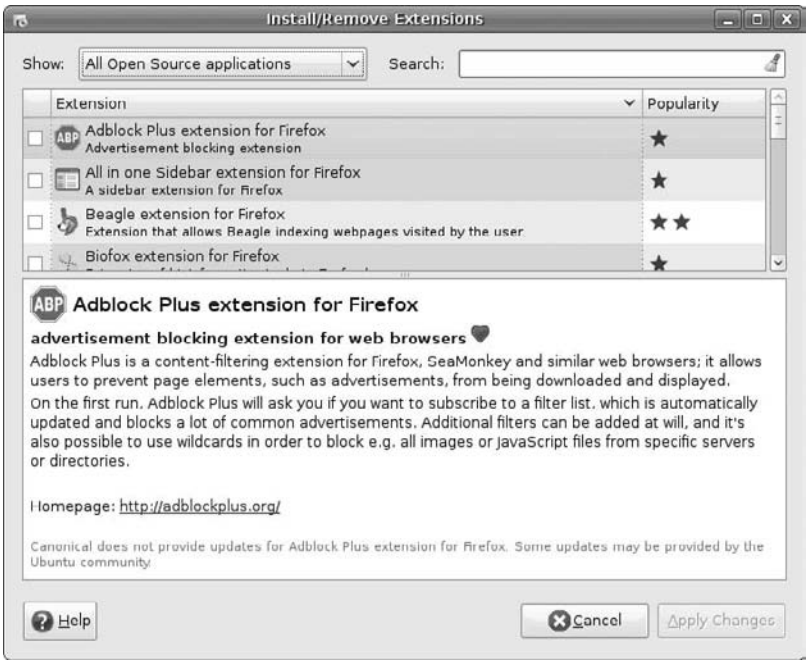
TIP **So many interesting extensions are available for Firefox that it's tempting to download and install a huge number of them at the same time, just to see which ones you like. I suggest that you try to control yourself and install only one or two at a time, for two reasons. First, some extensions may conflict with others, and this can cause problems with Firefox. Second, if you install too many extensions, it's hard to tell which one is responsible for which changes in Firefox. This can make it tough if you want to install the same set of extensions on multiple machines because you may waste time installing extensions that you don't actually want or need.**

An alternative to browsing the available add-ons at the Mozilla site is to use the handy Get Ubuntu Extensions link that is available on the Get Add-ons dialog shown in Figure 11-22. Clicking this link displays a dialog like the one shown in Figure 11-28, which provides a convenient way of quickly installing extensions that are tested and recommended by the Ubuntu folks. (Note that this dialog only enables you to install extensions, not themes.) To display a list of all recommended plug-ins, select All Open Source Applications from the Show: drop-down menu. If you are running the standard GNOME version of Ubuntu, you'll also note that this is the same dialog used by the Add/Remove command on the Applications menu, except that the name of the dialog is "Install/Remove Extensions," not "Add/Remove Programs."

NOTE **If you are running Kubuntu, you will see the Get Ubuntu Extensions link only if you have installed the ubufox package, as discussed in the section of this chapter entitled "Starting Firefox."**

FIGURE 11-28

Installing and removing recommended Ubuntu extensions



Some of these extensions are specific to Ubuntu (such as the Firefox Launchpad plug-in, if you're doing development or monitoring bugs, or the Ubuntu-it extension, if you're a member of the Italian Ubuntu community or speak Italian in general), while most are generally useful (such as AdBlock, for everyone, or the Web Developers toolbar, if you're a web developer). Clicking any extension in this list displays summary information about that extension in the bottom half of this pane.

To install any of the extensions in this list, click the checkbox to the left of its name in the top pane. The first time that you select any Open Source Application, you'll see a dialog warning you that you are about to install software that is maintained by the Ubuntu community, not Canonical. Well, you can hardly expect Canonical to maintain everything! Click Enable to accept that you want to do this. Once you're finished selecting extensions to install, click Apply Changes. This invokes the Synaptic package manager to install the extension, so you'll be prompted for your password. After Synaptic finishes installing the extension, the dialog shown in Figure 11-29 displays, asking if you want to install other Ubuntu extensions.

FIGURE 11-29

Prompt dialog to continue installing Ubuntu extensions



To install other extensions, click Add/Remove More Software. If you're done installing extensions from this dialog, click Close to return to the Add-ons dialog. Installing extensions using the Add/Remove Extensions dialog does not cause the Installation icon to display in the Add-ons dialog toolbar, but you will still have to restart Firefox to activate the newly installed extensions.

TIP **If you've installed new themes, you'll note that the last one that you installed is the new default theme used by Firefox. This will certainly cause both Firefox and subsequent dialogs to look different, as you'll see in some of the screenshots in the rest of this chapter. These use the Noia theme, which is one of my favorites and well worth checking out.**

After installing new extensions and themes and restarting Firefox, you can configure preferences for many of the extensions that you've installed. Some extensions, such as the AdBlock extension that helps block ads on Web pages, automatically display a Preferences dialog when they are first started so that you can supply mandatory configuration information, as shown in Figure 11-30.

FIGURE 11-30

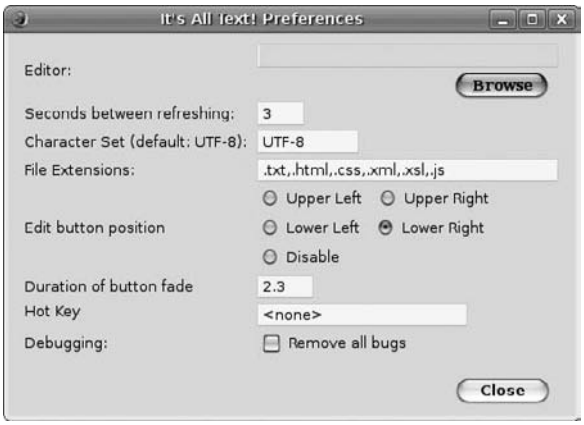
Setting preferences when an extension first starts



Most other extensions support some amount of configuration through preference settings. You can access the preferences dialogs for such extensions through entries for installed extensions on the Tools menu, or by searching for them in the extensions list and clicking the Preferences button in the entry for that extension. Figure 11-31 shows the preferences dialog for one of my favorite extensions, It's All Text!, which enables you to use the text editor of your choice (which should be Emacs) to edit text areas on Web pages.

FIGURE 11-31

Setting preferences for an extension through its Preferences dialog



This extension provides one of my favorite configuration items ever — the Remove all bugs checkbox! Sadly enough, this doesn't do exactly what it says, but instead activates verbose debugging so that the developer can diagnose and resolve problems. Maybe some day . . .

TIP **After experimenting with your new capabilities in Firefox, you can always uninstall an extension that you don't like or want by selecting the Add-ons command from the Tools menu, selecting the Extensions icon, highlighting the extension in question, and clicking Uninstall. The extension will be uninstalled the next time that you start Firefox.**

Adding New Firefox Plug-ins

As mentioned at the beginning of the "Enhancing Firefox" section of this chapter, plug-ins are applets that work within the Firefox framework and provide support for playing different types of multimedia files within your browser. The standard Ubuntu installation of Firefox comes with a number of plug-ins that support popular multimedia file formats such as Flash (SWF), Microsoft's Direct Video Express (DivX), RealPlayer, QuickTime, Windows Media Player, and so on.

NOTE **Unfortunately, many of the plug-ins provided by a standard Ubuntu installation are specific to GNOME multimedia players such as Totem (discussed in Chapter 19), which are not installed on a Kubuntu system. You will have to install Totem manually on your Kubuntu system if you want to be able to use the plug-ins that it provides.**

To see the plug-ins that are currently installed on your system, select the Add-ons command from the Tools menu, and click the Plugins icon. This displays the Plugins dialog, shown in Figure 11-32, which displays a list of all of the plug-ins that are currently installed on your system.

FIGURE 11-32

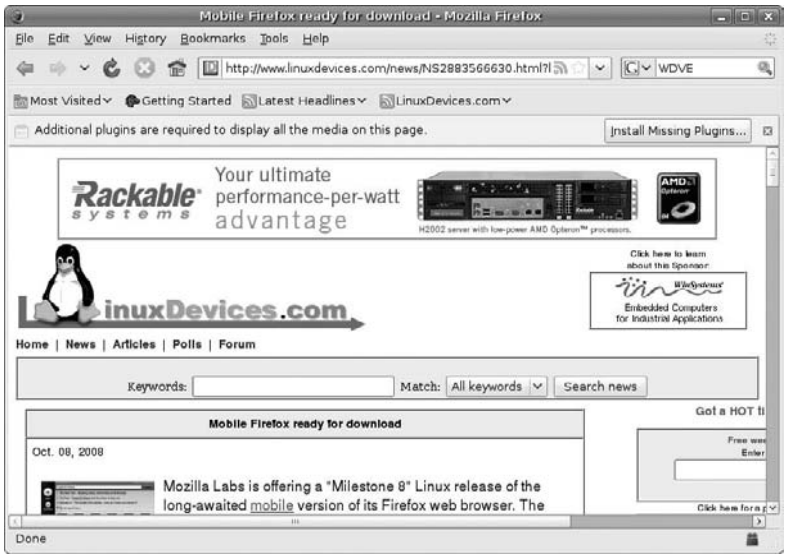
The Plugins dialog in Firefox



Generally, plug-ins are requested when you view a page that contains content which requires a plug-in to play it, as shown in Figure 11-33.

FIGURE 11-33

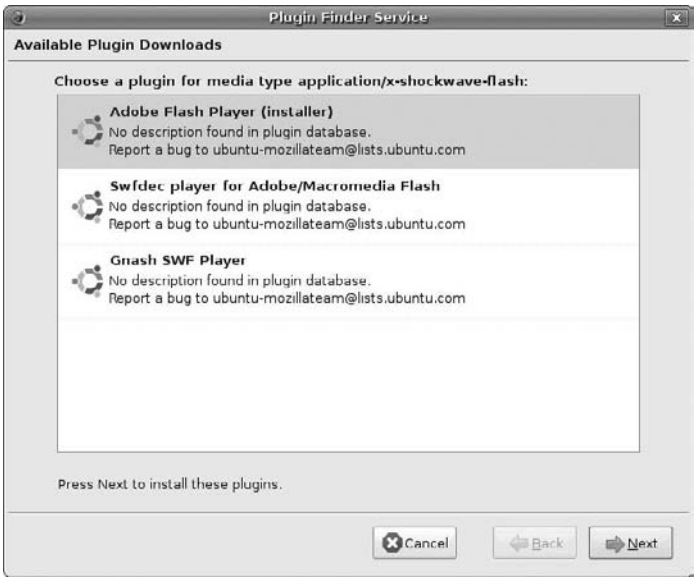
A Firefox page that requires a plug-in



When using Firefox on systems other than Ubuntu, clicking this button gives you access to the missing plug-in, enabling you to install it with a few mouse clicks (if you're lucky). You often have to install such plug-ins manually, depending on how well the developer of the plug-in designed and packaged it. Luckily, Ubuntu's ubufox package includes a plug-in finder service that is integrated with the GNOME Ubuntu package management tools, so clicking on the Install Missing Plugins button displays a page like the one shown in Figure 11-34, which displays a list of relevant plug-ins and enables you to select the one you want.

FIGURE 11-34

Ubuntu's plug-in finder service



After selecting the plug-in that you want, click Next. The Plugin Finder Service then uses the standard Ubuntu software installation/update mechanisms to install the plug-in. Nice! You will see a prompt dialog like the one shown in Figure 11-35 and, if necessary, you will be prompted for your password.

FIGURE 11-35

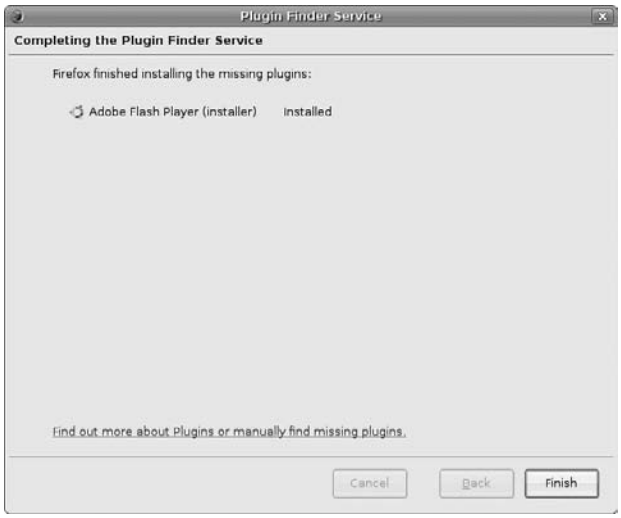
Confirming plug-in installation



Once the plug-in has been installed, you will see a dialog like the one shown in Figure 11-36, which enables you to exit from the Plugin Finder or open a browser window to the location on the Mozilla site where additional plug-ins are available (`https://addons.mozilla.org/en-US/firefox/browse/type:7`).

FIGURE 11-36

Confirming plug-in installation



You can click the link at the bottom of this dialog to browse for other plug-ins, or click Finish to close this dialog and return to Firefox.

Unlike extensions and themes, you can't remove plug-ins once they've been installed; you can only disable them. To disable an installed plug-in, click the Add-ons dialog's Plugins icon, and select the Disable button to the right of the entry for the plug-in that you want to disable. This enables you to install other plug-ins to handle a given type of content without losing track of the one(s) that you've already tried.

Summary

This chapter provided the highlights of configuring, using, and customizing Firefox, the most popular browser for Linux systems and one that is growing in popularity on other types of systems such as Microsoft Windows and Mac OS X. Firefox is the culmination of years of development and evolution in the Web browsing and open source communities. After using Firefox for a while on Ubuntu, you're sure to find that it's certainly worth installing and using on any other computer systems that you might have to use. You'll be happier because you'll have the same browser experience no matter what system you're using, and you're certainly not going to get any complaints from the open source community in general.

Chapter 12

Migrating from Windows Systems

Whether you're currently investigating Ubuntu as an alternative operating system to run on the computers in your home, or you're desperately looking for an alternative to operating system cost extortion, I hope that you'll eventually decide that switching to Ubuntu as your default operating system is the right thing to do. The political aspects of Linux are very important to many people (myself included), but most people simply want to be able to use their computers to get their work (or hobbies) done. Ubuntu does that. That it is stable, easy to upgrade, and virus-proof is just gravy.

Few people start out with Ubuntu as their first operating system and application environment. If you're moving to Ubuntu from another Linux distribution, you're in luck, because you can almost always find Ubuntu versions of the Linux applications that you're familiar with. However, if you're moving to Ubuntu from a computer that runs another operating system, switching to Linux is a bit more complex. After all, you've already configured your browser, e-mail, and networking. The computer that you're currently using is probably loaded with documents, spreadsheets, photographs, graphics, music, video files, and a lot of other data that you don't want to have to recreate (or simply may not be able to).

Moving existing files and personal data to your new Ubuntu or Kubuntu system is discussed in the last section of this chapter. Copying files from one computer to another is easy, but sometimes copying isn't enough. Some of the most important data that any of us have is our e-mail, favorite Web pages, and related configuration information. If you're coming from a traditional Microsoft Windows system, that data is locked into file formats and directory structures used by the applications that created it, such as Microsoft Outlook, Outlook Express, and Internet Explorer. Therefore, most of this chapter focuses on accessing your existing e-mail, browser favorites, cookies, and similar settings on your new Ubuntu system by converting that data into formats that are used by equivalent Ubuntu applications and then copying the converted data to your Ubuntu/Kubuntu system.

IN THIS CHAPTER

Basic migration scenarios

Backing up

Migrating e-mail

Migrating Web settings

Linux/Windows application equivalents

This chapter focuses on migrating your data from a Microsoft Windows system to a Ubuntu or Kubuntu Linux system. The same general concepts and tasks often apply if you're migrating from a Mac OS X system, but many of the details will differ. Although most of the examples in this chapter relate to different flavors of Windows, I will also provide tips that highlight Macintosh migration tasks without providing step-by-step instructions. I'll also use Ubuntu to generically refer to both Ubuntu and Kubuntu systems throughout this chapter — otherwise, we'd both have to see the phrase "Ubuntu and Kubuntu" a few hundred times.

Overview of Migration Scenarios

Chapters 2 and 3 discussed how to install Ubuntu, with a focus on installation on a new machine or on an existing machine where you can dedicate specific partitions to Ubuntu so that you can dual-boot. When migrating from a Windows system, you obviously need to be able to access the data from your Windows system. How you get to that information differs based on the type of Ubuntu installation that you're doing:

- **Installing Ubuntu on a new machine and leaving Windows on another:** This is the easiest migration scenario because you can copy file data from your Windows system to your Ubuntu system via a network connection or by burning CD or DVD copies of your Windows data. You can do any necessary file conversions on the Windows system, leaving your original data intact there, and then copy the converted data to your Ubuntu system to verify it.
- **Dual-booting an existing machine:** Still an easy scenario, you can mount your Windows partition(s) on your Ubuntu system and copy file data to Ubuntu. You may have to boot into Windows in order to convert some of your files, but you can copy the converted files directly from your Windows filesystem once you are running Ubuntu.
- **Converting a single system:** The trickiest scenario, you must use a backup of your Windows data as the location from which you can copy files. You must also do all necessary file conversions on the Windows system before you do the backup, or on the Linux system after you are running Ubuntu.

Both of the last two scenarios require making changes to your existing Windows system, so it is critical that you do a complete backup of your Windows system before installing Ubuntu or Kubuntu. See the next section for details.

This chapter discusses various ways of migrating your existing Windows data to an Ubuntu system. Frankly, I hope that you'll never need to read some of the sections of this chapter, but they're here if you need them. I suggest the following sequence of events in order to smoothly migrate your data to your new Ubuntu or Kubuntu system:

1. Start Outlook or Outlook Express and export your Address Book information in a generic format such as a CSV (comma separated value) file. You may not need this, but it's always handy to have a spare copy of this information in an application-independent format such as a CSV file.
2. If you are using Outlook, export your tasks and calendar information so that you can refer to it later. If you're using Outlook 2007, you can save your calendar in iCalendar format, which you can use with both Evolution and Kontact. Otherwise, you can at least save your tasks and calendar information in CSV format so that you can import that data in Evolution or Kontact.
3. If you are using Outlook, have created PST archives of your mail in the past, and want to migrate the contents of those archives to your new Ubuntu system, create a new folder and restore those

PST backups into it. (PST files are Microsoft Personal Folder Files used for backup purposes within Outlook.)

4. Install Thunderbird on your Windows system and allow it to import your existing mail and account information from Outlook or Outlook Express.
5. If you are using Outlook 2003 or later, create a new PST file in Windows 97–2002 format, and copy all of the mail folders from your existing mail folder there. Outlook 2003 introduced UTF-8 encoding in PST files, whereas previous versions of Outlook used ANSI-encoded PST files. I'm not aware of a Linux utility that can read UTF-8 PST files, but many utilities support the older format. You may not need this version of your existing PST file, but it's always handy to have a spare copy of this information in a standard format, just in case you ever need to access these files and don't have a Windows machine handy. See the section "Migrating Outlook Mail with readpst" later in this chapter, for detailed information about doing this.

NOTE **You can safely create this PST file after converting your mail with Thunderbird. Creating the new PST file before doing the conversion will cause Thunderbird to convert both of your PST files — not a big problem, but the duplication might be confusing.**

6. Install Firefox on your Windows or Mac OS X system and allow it to import your existing Favorites, saved passwords, cookies, and so on from Internet Explorer.
7. Export your Internet Explorer Favorites or Safari bookmarks to an HTML file. You may not need this, but it's always handy to have a spare copy of this information in an application-independent format such as HTML.

You will then want to back up all of the exported files that you created in the previous steps and your personal data (as described in the next section) from the Windows system to a CD, DVD, external drive, or network location, and verify that you can read the data that you've backed up on another computer system before proceeding.

Ready? Here we go . . .

Backing Up Your Personal Data

Backups are always a good idea. Few things are as frustrating as the realization that you've somehow lost years of e-mail, all your Web settings, the great American novel, all of your tax returns for the last decade, your music collection, and the digital photographs and video for all of your vacations and special occasions. I hope that you're already doing backups, but if you're installing Ubuntu on a machine that is currently running Windows, this may be your last chance to preserve all of the personal data on your computer. I've never encountered a problem when setting up a computer to dual-boot Windows and Ubuntu, but there's always a first time. Back up your data before proceeding! If you're converting a system that is currently running Windows to an Ubuntu system, your existing disk will be erased, so back up your data before proceeding!

NOTE **This is not a Windows book, and I cannot guess all of the locations where you or random Windows applications may have stored important data. This section discusses some standard locations for personal information. I am not responsible if you miss some critical files. Your mileage may vary. Caveat emptor! Caveat reader! Please make sure that your backups contain all of the files that you want to preserve before making any changes to your Windows system.**

When backing up your Windows system with an eye toward migrating existing information to your new Ubuntu system, you must create your backups in a format that is accessible from an Ubuntu system. The easiest way to do this is by simply copying files and directories to a CD or DVD. You will need to back up at least the following directory on your Windows system:

```
C:\Documents and Settings\Username
```

The following are the critical directories in this directory that contain your personal data:

- The Application Data and Local Settings folders where browsers and e-mail systems store your favorite Web sites, cookies, e-mail, account settings, passwords, and so on
- The My Documents folder, where your documents, music, pictures, and videos are usually stored

You should also back up any other directories from your local disk in which you have explicitly stored files that you've been working on. This can include project directories, personal directories that you've created on your C: drive, and so on.

When backing up your Windows system so that you can migrate existing data, do not use a PC backup program for which no Linux version exists. (That's pretty obvious, but it had to be said.) Frankly, even if you find a backup program that works on both Linux and Windows systems, I would not recommend using it. Backups that you can mount and access as standard ISO-9660 CDs and DVDs are easy to use and require no special software at any time. You can also easily test CD/DVD backups of files and directories on another system — do not assume that they will be readable from your new Ubuntu system without testing them. "Back up once, verify twice," as my grandfather used to say. You should verify that you can read the data that you've backed up on at least one other computer system before proceeding.

Migrating E-mail to Linux

Because e-mail is one of the two killer applications for modern computing (the other being the Web browser), preserving your existing e-mail is a pretty important aspect of moving to a new computer. This is especially true if you're also moving to a new operating system such as Ubuntu. I'm pretty compulsive about saving e-mail — I still have messages that I exchanged with friends in the early 1990s. (I used some proprietary and academic mail systems before that, but I'm not 100 percent compulsive.) You may not be as much of an electronic pack rat as I am, but you still probably want to preserve any mail messages that you've bothered saving.

Migration is necessary when moving from one operating system and local mail client to another because different clients store your account settings and e-mail in different formats. (Migration is not necessary if your mail is stored elsewhere and accessed through a Web client — this section applies only to e-mail that is stored on your computer.) While there are many different mail clients for Windows systems, most people use either Microsoft Outlook Express, a free mail client that is typically included with Windows Internet Explorer downloads, or Microsoft Outlook, a commercial sibling that is part of the Microsoft Office Suite. The biggest thing that these two applications have in common is that they are both often described as virus propagation software disguised as an e-mail client, but their primary differences are core capabilities and the format in which they store mail. The DBX format used by Outlook Express is easy to migrate, as described in the next two sections. Migrating Outlook mail from PST files requires some special handling, which I'll discuss later in the section "Migrating Outlook Mail with readpst." Most Linux mail clients use a format known as mbox (mailbox), which has been used by some Windows

mailers (such as Eudora), and was also used by Mac OS X's Mail.app up to version 2.0, after which Apple switched to a proprietary mailbox format.

NOTE **Migrating your e-mail is necessary only if your e-mail account is a POP/POP3 mail account or you are using an IMAP account with Microsoft Outlook and have archived your mail on your local system. If your e-mail account is a standard IMAP account and you are either not using Microsoft Outlook or don't have any local archives, all of your saved mail is still on the IMAP server, and no migration is necessary.**

The following is a quick overview of the pros and cons of the Outlook Express migration mechanisms discussed in the next two sections.

- **Migration via Thunderbird:** Convenient because your mail files and account settings are converted automatically. However, you have to download and install Thunderbird on your Windows system and do the conversion before you do your backups. You also have to search for the converted files when copying them to your Ubuntu system or reading them from your backups. Unless you continue to use Thunderbird on Ubuntu, you'll also have to import the converted files into Evolution or Kontact format, one converted file at a time.
- **Manual migration:** Convenient because you can selectively convert individual mail files. However, if you are migrating Outlook Express mail, note that the readdbx application isn't provided on Ubuntu systems by default, so you have to compile it yourself, which also means that you have to install the GNU C compiler on your Ubuntu system. See Bonus Chapter 3, "Software Development on Ubuntu," for information about compiling applications. If you are migrating Outlook mail, the readpst utility is available in the Ubuntu repositories, but it can only read the PST files used by versions of Outlook prior to Outlook 2003. You'll have to create a PST file in the old format and then copy the mail folders that you want to migrate into that file. Regardless of whether you're migrating Outlook Express or Outlook mail, these utilities run only on your Ubuntu system, so you will then have to search for the files that you want to convert on your backups and convert them one at a time. Finally, you'll have to import the converted files into Evolution or KMail, one converted file at a time.

I suggest that you use Thunderbird for the migration. It's really much easier. However, if something goes wrong and you need to migrate e-mail from your backups when you no longer have a Windows system around, the readdbx and readpst mechanisms that are discussed later in this chapter can be lifesavers.

Migrating Mail via Thunderbird

Thunderbird is a popular mail client from the Mozilla Foundation (`www.mozilla.com/thunderbird`) that is available for the Microsoft Windows, Apple Macintosh, and Linux platforms. Like Firefox, Thunderbird supports plugins that make it easy to customize and extend its default behavior.

Aside from its excellence as a standalone mail client, Thunderbird can be very handy when migrating e-mail to Ubuntu from a Windows system. You can use Thunderbird to migrate your Outlook Express or Outlook mail in one of two ways:

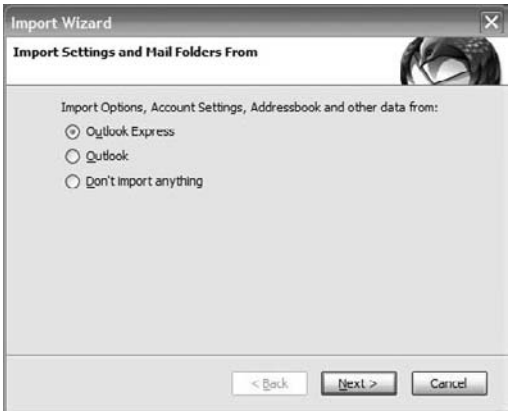
- If you're already using Thunderbird on your Windows system and want to continue using Thunderbird on Ubuntu, you can install Thunderbird on Ubuntu, copy your existing Thunderbird profile to your new Ubuntu system, and everything will "just work."
- If you are using Outlook Express on your Windows system, you can install Thunderbird on your Windows system and let it convert your Outlook Express e-mail and account settings for you.

As you might expect, you must do these migrations on your Windows system, and you should probably do them before doing your final backups, so that both your original mail files and the converted versions are included on your backups.

You can download a copy of Thunderbird for Windows by clicking on the download link at `www.mozilla.com/thunderbird`. The last installation dialog offers to automatically start Thunderbird, which you should do. This displays the dialog shown in Figure 12-1, asking if you want to import your Outlook Express mail and account settings into Thunderbird.

FIGURE 12-1

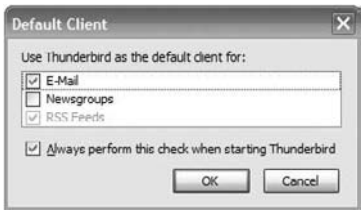
Importing Outlook Express mail into Thunderbird



The Outlook Express radio button is selected by default — select the Outlook button if you want to migrate mail from Outlook instead. (You'll see the Outlook button only if Outlook is installed on your computer.) Click Next to proceed to a status dialog that reports the progress of importing your mail account settings, address book, and mail folders into Thunderbird. After displaying a final status dialog, Thunderbird will start on your system, displaying a dialog like the one shown in Figure 12-2.

FIGURE 12-2

The Default Client dialog in Thunderbird



Although you're migrating all of your mail to Ubuntu and will hopefully never need to read mail on Windows again, you may want to click OK with the E-Mail checkbox selected. This will modify your Windows mail application settings so that Windows will start Thunderbird, rather than Outlook or Outlook Express, if you accidentally click a link that causes you to read mail on your Windows system.

NOTE **Unless you are using IMAP, you should not read your mail on your Windows system using Outlook or Outlook Express after migrating your mail to Thunderbird. If your e-mail account is a POP account and you read mail, new mail will be downloaded to your Outlook or Outlook Express mail files, but will not be reflected in the Thunderbird mail files that you converted. Subsequently reading mail with Thunderbird on your Windows system is OK because your mail will be stored in the Thunderbird mail files that you're migrating.**

At this point, you should create an external backup of at least the directories discussed in the section "Backing Up Your Personal Data."

NOTE **After all this talk about Thunderbird, you're probably wondering "Why the heck didn't Bill include a chapter on Thunderbird in this book?" I gave that serious thought, but eventually decided that encouraging people to use the mail clients provided in the Evolution and Kontact PIMs was the right thing to do. Thunderbird is an excellent mail client — many of my friends use and swear by Thunderbird, and I have used it on a laptop or two in the past. However, I think that most Ubuntu and Kubuntu users will get much more long-term value out of sticking with Evolution or Kontact.**

Where Thunderbird Stores Its Mail

Thunderbird stores the mail for each user in the user's profile directory. On a Windows system, your profile directory is located in the following directory, where *Username* is your login name on the Windows machine (sometimes concatenated with the hostname of your Windows system):

```
C:\Documents and Settings\Username\Application Data\Thunderbird\Profiles
```

You should be able to navigate to this folder using Windows Explorer. Thunderbird profile directories in this directory have names of the form *xxxxxxxx*`.default`. This directory contains all of your Thunderbird information, including a Mail subdirectory that contains all of your account information and mail files.

On Linux systems, mail is stored in a directory whose name has the same format, but its location depends on the version of Thunderbird that you're using:

- If you have downloaded, installed, and are using the official version of Thunderbird from the Mozilla project rather than using Synaptic or the Adept Manager, your profile directory and profile initialization file (`profiles.ini`) will be located in a directory called `.thunderbird` that is located in your home directory.
- If you have installed Thunderbird using Synaptic or the Adept Manager, your profile directory and profile initialization file (`profiles.ini`) will be located in a directory called `.mozilla-thunderbird` that is located in your home directory. This was an older name for the directory that holds Thunderbird initialization files, and has been preserved by the Ubuntu and Debian folks to guarantee compatibility with older Thunderbird installations.

In my opinion, it's unfortunate that Ubuntu and Debian have preserved the old name for the directory that holds Thunderbird initialization and profile information, but no one asked me.

TIP **On Mac OS X systems, your Thunderbird profile directory is stored in one of two locations, either** `~/Library/Thunderbird/Profiles` **or** `~/Library/Application Support/Thunderbird/Profiles`**.**

Installing or Updating Thunderbird

To migrate to a Ubuntu system on which you plan to install and use Thunderbird, you'll need to install the Ubuntu version of Thunderbird, which you can find as the thunderbird package using Synaptic or the Adept Manager. The related mozilla-thunderbird package is a compatibility package for users who are upgrading older systems, and simply requires the Thunderbird package.

How to proceed depends on whether you've just installed Thunderbird for the first time on your Ubuntu system, or whether you're already been using Thunderbird on your Ubuntu system for a while:

- If you have just installed Thunderbird for the first time, proceed to the next section ("Importing Mail into Thunderbird: New Users") for information on what files to copy from your Windows system, where to put them on your Ubuntu/Kubuntu system, and how you'll need to modify the Thunderbird profile file.
- If you've installed and used Thunderbird on your Ubuntu or Kubuntu system before, you'll want to preserve your old Thunderbird mail, and should therefore skip to the section "Importing Mail into Thunderbird: Existing Users."

Importing Mail into Thunderbird: New Users

If you have never used Thunderbird on your Ubuntu system before, you can copy over the converted mail and profile initialization information from your Windows system and have Thunderbird use it immediately.

To do this, first create the directory `~/.mozilla-thunderbird` before starting Thunderbird. Next, recursively copy the `xxxxxx.default` directory from your Windows backups to the `~/.mozilla-thunderbird` directory that you just created in your Ubuntu home directory. You can do this by copying this directory from the CD, DVD, or network location to which you backed up the following directory:

```
C:\Documents and Settings\Username\Application Data\Thunderbird\Profiles
```

TIP **If you are installing Ubuntu on a new machine, you can use Samba to mount a shared Windows drive on your Ubuntu system, and copy the files over the network. See Bonus Chapter 1 for information on mounting shared Windows drives on your Ubuntu system.**

You should also copy the following file from the backup of this directory from your Windows system into your `~/.mozilla-thunderbird` directory:

```
C:\Documents and Settings\Username\Application Data\Thunderbird\profiles.ini
```

When copied from a Windows system, this file will initially look something like the following:

```
[General]
StartWithLastProfile=1
[Profile0]
Name=default
IsRelative=1
Path=Profiles/hus5hy55.default
```

After copying it to your Ubuntu system, use a text editor to modify the version of this file that you copied from your Windows system and remove the Profiles directory from the Path variable. An updated version of this same file would look like the following:

```
[General]
StartWithLastProfile=1
```

```
[Profile0]
Name=default
IsRelative=1
Path=hus5hy55.default
```

When you start Thunderbird for the first time on Ubuntu, it will instantly use the profile directory from your Windows system, and things should look exactly the same as they did there. If this is the case, congratulations! Your mail has been migrated successfully, and you can now resume your normal e-mail existence using Thunderbird as your default mail client.

TIP **If you subsequently want to switch to using Evolution or Kontact, you can always import your Thunderbird mail into those applications as described later in the sections "Importing Thunderbird Mail into Evolution" and "Importing Thunderbird Mail into Kontact."**

Importing Mail into Thunderbird: Existing Users

If you've already been using Thunderbird on Ubuntu and are simply migrating old mail from a Windows system, the directory `~/.mozilla-thunderbird` will already exist in your home directory on your Ubuntu system. That directory will already contain a default profile directory named something like *`xxxxxx`*`.default` and your profile initialization file, `profiles.ini`.

In this case, you will want to use Thunderbird's import capability to import the Thunderbird mail from your Windows system into your existing Thunderbird mail files. You will therefore need access to the directory where Thunderbird stored the converted versions of your Windows mail. This should be located on the CD, DVD, external drive, or network location to which you backed up the `C:\Documents and Settings\Username\Application Data\Thunderbird\Profiles` directory.

Your Thunderbird mail is actually stored in a subdirectory of this directory, with a name like *`xxxxxx`*`.default/Mail/Local Folders`, where *`xxxxxx`* is a random string that was generated when you first started Thunderbird. This directory contains files that contain your mail and auxiliary files and directories used to store Thunderbird configuration information. To determine which files you need to import into Thunderbird, start an Ubuntu xterm or console window, change directory to the *`xxxxxx`*`.default/Mail/Local Folders` directory, and execute a command like the following:

```
$ find -not -empty | egrep -v '.msf$|.sbd$|.dat$|^.$'
./Personal Folders.sbd/Inbox
./Personal Folders.sbd/Personal
./Personal Folders.sbd/Writing
```

In this case, the files Inbox, Personal, and Writing are all non-zero in size (i.e., not empty) and thus contain mail that you will want to import. The egrep portion of the command filters out (in order) Thunderbird Mail Summary Files (MSF), empty Mail Subdirectories (SBD), and the DAT files in which Thunderbird stores registry and version data.

Once you've identified the files that contain the Windows mail that you want to import into your existing Thunderbird mail, you will need to copy these files to the directory in your existing Thunderbird profile so that you can access them from Thunderbird. You can do this in one of two ways:

- If you do not have Thunderbird message folders with the same name on your Ubuntu system, you can simply copy them with their existing names.
- If you already have Thunderbird message folders with the same names on your Ubuntu system, you will have to copy your Windows message folder and give them new names so that they don't overwrite the existing Thunderbird message folders with the same name.

For safety's sake, I always use the second approach because it's easy to overwrite folders with important names such as Inbox. After changing directory to the *xxxxxx*.default/Mail/Local Folders directory on your Windows backups, commands like the following would copy the mail folders that we previously identified to your Ubuntu installation, adding the -imported suffix to the copied folders:

```
cp "./Personal Folders.sbd/Inbox" \
   ~/.mozilla-thunderbird/hf8dwpmj.default/Mail/Local\ Folders/Inbox-
imported
./Personal Folders.sbd/Personal \
   ~/.mozilla-thunderbird/hf8dwpmj.default/Mail/Local\ Folders/
Personal-imported
./Personal Folders.sbd/Writing \
   ~/.mozilla-thunderbird/hf8dwpmj.default/Mail/Local\ Folders/Writing-
imported
```

That's extremely tedious to type for any more than one mailbox, so the following is a script that requires the -t argument to specify that you are migrating Thunderbird mail, takes the name of your profile directory as an argument, and does this for you:

```
#!/bin/bash

function usage {
    echo "Usage: import_thunderbird_mail.sh [-e] [-t profile-dir]"
    echo "  -e : import Thunderbird mail into Evolution"
    echo "  -t : import Thunderbird mail into Thunderbird"
    exit 1
}

if [ $# == 0 ] ; then
    usage
    exit
fi

while getopts ":t:e" opt
do
  case $opt in
    e  )
        TARGETDIR=~/.evolution/mail/local
        ;;
    t  )
        if [ $# != 1 ] ; then
          echo "No Thunderbird profile directory specified."
          exit 1
        fi
        TARGETDIR=~/.mozilla-thunderbird/$1/Mail/Local\ Folders
        ;;
      \? ) usage
        exit 1;;
  esac
done
shift $(($OPTIND - 1))
```

```
find . -not -empty | egrep -v '.msf$|.sbd$|.dat$|^.$' | \
    while read mbox ; do
      newmbox=`echo $mbox | sed -e 's;.*/;;' -e 's;$;-imported;'`
      echo "Copying $mbox to $newmbox"
      cp "$mbox" $TARGETDIR/$newmbox
    done
```

NOTE **You can download this script, and others from this book, from my Web site at** `www.vonhagen.org/Ubuntu/scripts`. **The name of this script is** `import_thunderbird_mail.sh`.

You can find out the name of the Thunderbird profile directory by listing the directory `~/.mozilla-thunderbird`:

```
$ ls ~/.mozilla_thunderbird
hf8dwpmj.default   profiles.ini
```

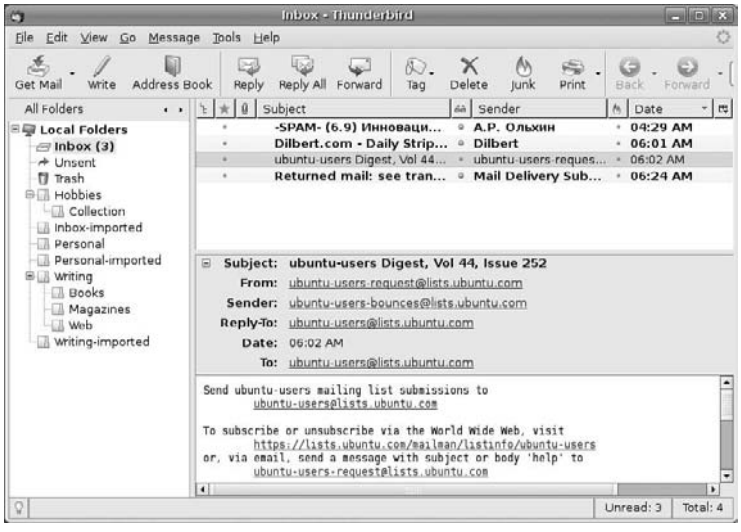
In this case, your profile directory is the directory `hf8dwpmj.default`. In this case, the command to run this script and the output from its execution would be the following:

```
$ import_thunderbird_mail.sh -t hf8dwpmj.default
Copying ./Personal Folders.sbd/Writing to Writing-imported
Copying ./Personal Folders.sbd/Inbox to Inbox-imported
Copying ./Personal Folders.sbd/Personal to Personal-imported
```

The next time that you start Thunderbird on Ubuntu, you should see the imported mailboxes, as shown in Figure 12-3.

FIGURE 12-3

Existing and imported mailboxes in Thunderbird



Congratulations! Your mail has been migrated successfully, and you can now copy messages from the imported folders into the folders that you are currently using in Thunderbird. You can also simply rename these folders, removing the `-imported` extension, if you want to preserve the familiar folder names that you used when you were a Windows user.

Importing Thunderbird Mail into Evolution

Letting Thunderbird convert your existing Outlook/Outlook Express e-mail (as described in the section "Migrating Mail via Thunderbird") is an excellent first step for importing your Windows mail into Evolution, the GNOME PIM and mail client used by the standard Ubuntu desktop distribution. Both Thunderbird and Evolution store the mail for a mail folder in a file in the traditional UNIX mbox (mail box) format. This makes it extremely easy to import mail from Thunderbird into Evolution.

Evolution provides an import wizard that imports all mail from older Linux mailers (Netscape Navigator, and so on), and also lets you import single files in supported formats such as Berkeley Mailbox (mbox), Evolution CSV or Tab (.csv or .tab), Mozilla CSV or Tab (.csv or .tab), Outlook CSV or Tab (.csv or .tab), Vcard (.vcf or .gcrd), LDAP Data Interchange Format (.ldif), vCalendar (.vcf), and iCalendar (.icf). Using Evolution's import wizard to import single files is discussed later in this section, and makes sense if you are importing one or two files. However, single-file import can be tedious and time-consuming if you have a large number of mailbox files that you want to import. In the latter case, you can use the `import_thunderbird_mail.sh` script that was discussed in the previous section to import all of your mailboxes into Evolution, adding a `-imported` suffix to each of their names so that you can identify the imported mailboxes.

NOTE **You can download this script, and others in this book, from my Web site at `www.vonhagen.org/Ubuntu/scripts`. The name of this script is `import_thunderbird_mail.sh`.**

To run this script and import mailboxes that you created by importing your Outlook/Outlook Express mail into Thunderbird on a Windows system, first change directory to the CD, DVD, or network location to which you backed up the `C:\Documents and Settings\`*`Username`*`\Application Data\Thunderbird\Profiles` directory. Your Thunderbird mail is actually stored in a subdirectory of this directory, with a name like *`xxxxxx`*`.default/Mail/Local Folders`, where *`xxxxxx`* is a random string that was generated when you first started Thunderbird. This directory contains files that contain your mail and auxiliary files and directories used to store Thunderbird configuration information. Change directory to the *`xxxxxx`*`.default/Mail/Local Folders` directory and run the `import_thunderbird_mail.sh` script with the `-e` argument to state that you are migrating mail into Evolution.

The command to run this script and the output from its execution is as follows:

```
$ import_thunderbird_mail.sh -e
Copying ./Personal Folders.sbd/Inbox to Inbox-imported
Copying ./Personal Folders.sbd/Personal to Personal-imported
Copying ./Personal Folders.sbd/Writing to Writing-imported
```

The next time that you start Evolution, you will see the newly imported files, as shown in Figure 12-4.

As mentioned previously, you can also import selected mailboxes into Evolution, which may be more convenient if you're uncomfortable at the command line or simply don't want to have to download and install my shell script. To use Evolution's import wizard, start Evolution and select the File menu's Import option. A welcome dialog displays. Click Forward to proceed to the dialog shown in Figure 12-5, which enables you to identify whether you are importing data and settings from an older program, or whether you are importing a single file.

FIGURE 12-4

Existing and Imported mailboxes in Evolution

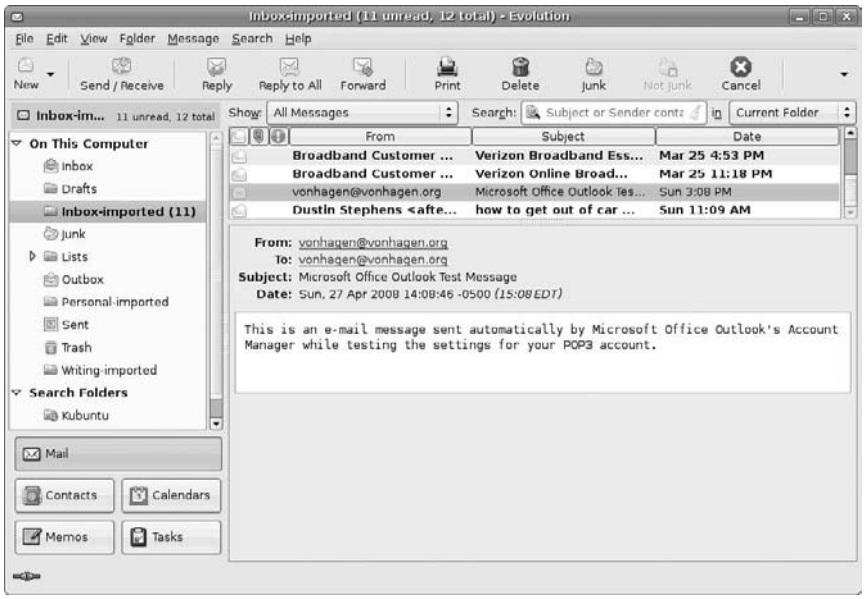


FIGURE 12-5

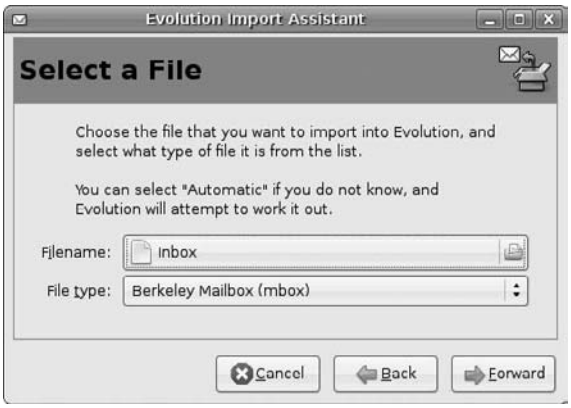
Identifying import type in Evolution's import wizard



Select Import a single file and click Forward to proceed. The dialog shown in Figure 12-6 displays.

FIGURE 12-6

Locating the file to import in Evolution's import wizard

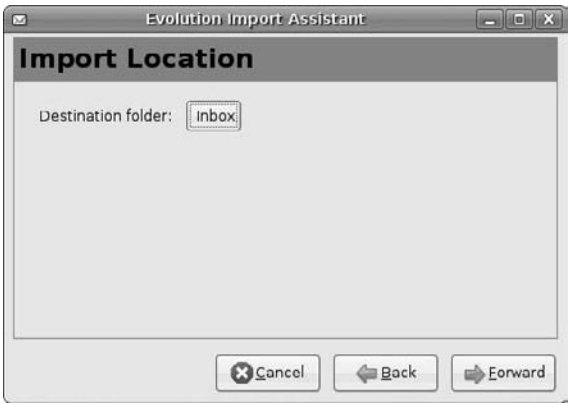


Click the Filename button to display a dialog that enables you to navigate to and select the file that you want to import. The mail files that you converted using Thunderbird should be located on the CD, DVD, external drive, or network location to which you backed up the `C:\Documents and Settings\Username\Application Data\Thunderbird\Profiles` directory. They will be located in a subdirectory of this directory that has a name like *xxxxxx*`.default/Mail/Local Folders`, where *xxxxxx* is a random string that was generated when you first started Thunderbird.

Click Open to close the navigation dialog and open the file that you want to import. The File type entry on the "Import a File" dialog should display Berkeley Mailbox (mbox) because this is the format of mail files that have been created by Thunderbird. (It is also the format of mail files used by Evolution.) Click Forward to proceed. The dialog shown in Figure 12-7 displays.

FIGURE 12-7

Specifying the Import Location in Evolution's import wizard



This dialog enables you to select the Evolution mail folder into which you want to import the messages in the mail file that you selected in the previous step. The default import location is your Inbox folder, but you can click the Destination folder button to navigate to and choose another folder into which you want to import your messages.

TIP **When importing mail for others, I sometimes create a new folder in Evolution named IMPORT, and use this as the folder into which I import messages. This makes it easy to differentiate between imported messages and messages that may already exist in a given folder. After finishing the import process and moving imported messages to other folders, I delete the IMPORT folder, and everything looks just as it did when I started (except for the number of messages in various folders).**

After selecting a folder, click Forward to proceed to the next dialog. A confirmation dialog displays. To proceed, click Import. Evolution displays a status dialog as it imports the messages. When this closes, congratulations! You've successfully imported mail from Thunderbird. You can now re-select the File Import command to import another mail folder, or begin filing and organizing the newly imported messages.

Importing Thunderbird Mail into Kontact

Kontact, discussed in Chapter 10, is the Personal Information Management application used on most KDE-based systems, including Kubuntu. Kontact is "komposed" of a number of formerly standalone applications that have been unified into a single PIM, which includes an excellent mail client (known as KMail when it was a standalone application).

Kontact provides an excellent import mechanism that not only imports mail in various formats from a huge number of mail clients, but also thoughtfully organizes your imported mail into a single hierarchical folder in the Kontact mail client. This makes it easy to see exactly which folders and messages were imported so that you can make sure that you move all of your historical mail to the right folders in Kontact.

Importing mail into Kontact as described in this section depends on your having imported your Outlook/Outlook Express mail into Thunderbird, as described in the section "Migrating Mail via Thunderbird" earlier in this chapter. Importing mail also depends on whether you backed up that data to a CD, DVD, external drive, or network location that you can access from your Kubuntu system (also described earlier, in the section "Backing Up Your Personal Data").

After starting Kontact on your Kubuntu system, make sure that you are in the e-mail component by clicking the envelope icon in the left-hand icon bar. Select the File menu's Import Messages command. The dialog shown in Figure 12-8 displays.

The drop-down menu in the center-right portion of the screen enables you to identify the application that created the mail messages that you want to import. At the time of this writing, the following are the supported applications from which mail can be imported:

- Evolution 1.*x*
- Evolution 2.*x*
- Lotus Notes
- OS X Mail
- Opera
- Outlook Express
- Pegasus Mail

- Plain Text
- Sylpheed
- The Bat!
- Thunderbird/Mozilla
- UNIX mailers using mbox format

FIGURE 12-8

The Import Messages dialog in Kontact



This list is here for your reference, but for the purposes of this book, we will only be concerned with the ability to import mail created by Thunderbird. I realize that this list shows that Outlook Express is a supported format, but going through Thunderbird is safer if you ever need to restore the same mail on another system. I also wanted to provide a single set of instructions for exporting mail from both Outlook and Outlook Express.

Select Import Thunderbird/Mozilla Local Mails and Folder Structure and click Next to proceed. After selecting this import format, the information pane on the import dialog changes to provide information about the location from which you should import mail, as shown in Figure 12-9.

FIGURE 12-9

Identifying the location of mail to import



After reading this information, click Next to display a navigation dialog that enables you to select the location where the mail that you converted using Thunderbird is stored. The mail files that you converted using Thunderbird should be located on the CD, DVD, or network location to which you backed up the `C:\Documents and Settings\Username\Application Data\Thunderbird\Profiles` directory. They will be located in a subdirectory of this directory that has a name like *`xxxxxx`*`.default/Mail/Local Folders`, where *`xxxxxx`* is a random string that was generated when you first started Thunderbird.

After selecting this directory, click OK to close the navigation dialog and begin the import process. A progress dialog displays as your mail is imported, as shown in Figure 12-10.

Once the import process completes, click Finish to close the Import messages dialog and return to Kontact. As shown in Figure 12-11, you will see a Thunderbird-Import message folder in your local folders list, which contains subdirectories and mail folders that reflect how your mail was organized on your Windows system.

Congratulations! You've successfully imported mail from Thunderbird into Kontact. You can now rest on your laurels or begin filing and organizing the newly imported messages.

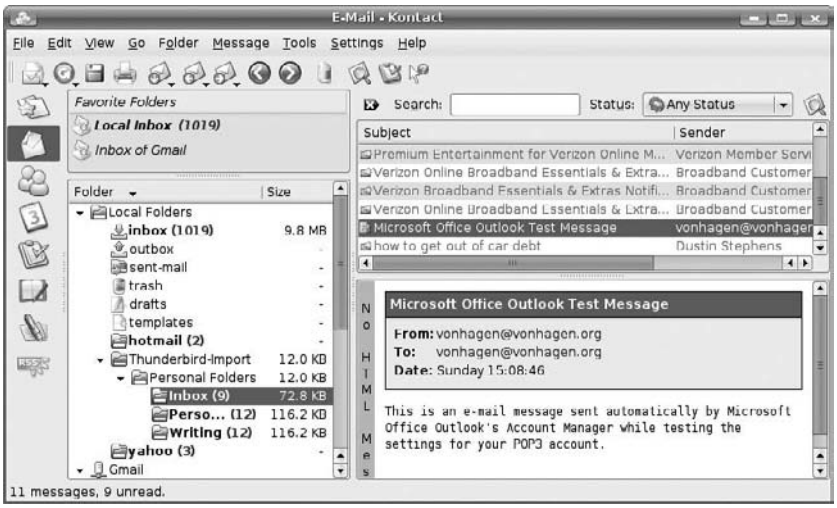
FIGURE 12-10

Kontact's import status dialog



FIGURE 12-11

Imported mail in Kontact



Manually Migrating Mail

The essentially automatic export mechanism that is provided by installing Thunderbird on a Windows system is extremely convenient and is also fairly bullet-proof because it uses Windows function calls that provide direct access to your Outlook Express mail. After importing your mail into Thunderbird, you can back up the directories in which Thunderbird stores its mail and use those to import your mail into Ubuntu and Kubuntu mail clients such as Evolution and Kontact. If you're lucky, you'll never need to see or use those backups again.

Unfortunately, not everyone is lucky. I've worked with clients who did their backups before using Thunderbird to generate portable versions of their mail files, managed to avoid exporting specific Outlook Express mail files into Thunderbird, or needed to restore mail from older Windows backups and import those restored folders. I don't even know how people manage to do some of these things, but these are all cases in which you can't use the traditional Thunderbird-to-Whatever mail migration method.

Luckily for us, the open source community is large enough that somebody has had to do these sorts of things before, and therefore have written tools to do these things. Migrating mail folders manually requires a bit more work than clicking a few buttons and letting Thunderbird do the heavy lifting, but it works just as well.

The next two sections describe how to use the command-line readdbx and readpst utilities to manually create mbox files from your DBX files (outlook Express mail) and PST files (Outlook mail). I hope that you won't ever need this information, but better safe than really, really sorry.

Migrating Outlook Express Mail with readdbx

Outlook Express stores its mail folders in files that are in DBX format, a database format of sorts. Luckily, the libdbx package provides two utilities to help convert files in these formats to files in the standard mbox format. These are the readdbx utility, which converts single DBX files, and the readoe utility, which will take a directory of DBX files (such as the directory where Outlook Express stores your mail) and will convert all of the files in this directory into mbox-format files in another directory.

The libdbx project is actually a part of the ol2mbox (Outlook to mbox) project on SourceForge (`http://sourceforge.net/projects/ol2mbox`). There hasn't been a new release of libdbx for a while because, well, the current version "just works." (The libpst package for reading Outlook mail from PST files is a different story, as discussed in the next section.) For all of us Ubuntu users, the bad news is that the libdbx package is not available in the Ubuntu repositories, so you'll have to download the source code and build it yourself if you need to use it.

I won't go into the details of building and installing the package here — Bonus Chapter 3 explains how to install the tools necessary to compile source code that is written in the C programming language. Here's a quick overview:

1. Download the source code package from the SourceForge project Web site by clicking the download link and retrieving the `libdbx_1.0.3.tgz` compressed archive file.
2. Uncompress the archive using a command such as `tar zxvf` *`filename`*, where *`filename`* is the name of the file that you downloaded.
3. Change directory to the directory to which the software was extracted (probably `libdbx_1.0.3`), and configure the software using the command `./configure`.
4. Use the make program to build and install the readdbx and readoe utilities by executing the command `make install`.

Once you've done all of these steps, you should be able to change directory to the directory that contains the DBX files that you want to convert. If you're working from backups of a Windows system, the "right" directory can be tricky to identify because it depends on the name of an internal identity that is created by Windows. The generic format of the name of this directory (located under your user directory) is as follows:

```
Local Settings\Application Data\Identities\ID\Microsoft\Outlook Express
```

In this example, you would change *ID* to some string that is your user identifier. You should be able to determine the "right" `ID` by exploring your backups.

To convert an entire directory of Outlook Express files to mbox format, the readoe utility must be run from the parent of this directory using commands such as the following:

```
cd ..
mkdir converted
readoe -i Outlook\ Express -o converted
```

All of the Outlook Express mailbox files in the Outlook Express directory will be converted to mbox-format files that will be created in the "converted" directory.

If you want to convert only a single DBX file (such as one that you've restored from backups), you'd execute a command like the following from the directory that contains your Outlook Express mail:

```
$ readdbx -f Personal.dbx -o personal.mbox
```

This command converts the contents of your Personal folder to the file `personal.mbox`. Once you've converted a mailbox, you can use your mail client's import capability to import the mbox file that you've created. Importing mail from mbox files was explained in the previous sections.

Migrating Outlook Mail with readpst

Microsoft Outlook stores its mail folders in files that are in PST format. PST stands for "Personal Storage Table," but are most often referred to as "Personal File Folder" files. The readpst utility, which is available in the Ubuntu repositories, enables you to read PST files, extract the mail folders that they contain, and create equivalent files in mbox or KMail format.

The libpst project was originally a part of the ol2mbox (Outlook to mbox) project on SourceForge (`http://sourceforge.net/projects/ol2mbox`), just as the libdbx project is. Recently, the need for changes to libpst has led to a new maintainer and home for libpst. You can still get an old version of libpst from the ol2mbox site, but you must get the latest versions of libpst from `http://alioth.debian.org/projects/libpst/`. For all of us Ubuntu users, the good news is that the readpst is also available in the Ubuntu repositories, so you can just install it using Synaptic or the Adept Manager — no compilation required!

One twist in using readpst to manually convert Outlook mail is that, beginning with Outlook 2003, Microsoft's PST format has changed. Earlier versions of Outlook used the ANSI character set in their PST files, while Outlook 2003 and later use the Unicode character set. Therefore, in order to extract mail from a new-format PST file, you should create an old-style PST file in Outlook 2003 and later on your Windows system. You can do this by using the File ➪ New ➪ Outlook Data File command, selecting Outlook 97-2002 Personal Folders File (.pst) as the type of file that you want to create. You can then right-click on the folders containing the mail that you want to be able to export, select the Copy command, and select the section in Outlook that is associated with the new PST file.

Once you've copied your mail folders to an old-format PST file, you should be able to change directory to the directory that contains the PST file that you want to convert. If you're working from backups of a Windows system, the name of this directory (located under your user directory) is as follows:

```
Local Settings\Application Data\Microsoft\Windows\Outlook
```

Once you have located an old-style PST file, create a directory to hold the converted mbox-format files, and execute the readpst command. An example command and related output for readpst looks like the following (using a directory called mbox as the location for the converted mbox-format files):

```
$ readpst -o mbox Personal\ Folders-97-2002.pst
Opening PST file and indexes...
About to start processing first record...
Processing items...
Processing Folder "Deleted Items"
        No items to process in folder "Deleted Items", should have been 0
Processing Folder "Inbox"
        "Inbox" - 12 items done, skipped 2, should have been 12
        "Personal Folders" - 0 items done, skipped 0, should have been 0
Finished.
```

The readpst utility will generate an error message and exit if you attempt to convert a Unicode PST file produced by Outlook 2003 or later.

You can then import the converted files directly from the mbox directory into Evolution or Kontact.

Migrating Web/Browser Settings to Linux

Chapter 11 of this book discussed Firefox, the default Web browser provided on standard Ubuntu desktop systems and a worthy addition to any Kubuntu system. (The default browser on Kubuntu systems is Konqueror, which was discussed in Chapter 7.) As a cross-platform browser that is officially supported on Linux, Windows, and Mac OS X systems, and available for Solaris, AIX, and other commercial UNIX distributions, Firefox is also a tremendous time-saver for when you have to use multiple operating systems but would like to standardize on one browser.

NOTE **This section focuses on migrating Web settings to Firefox because I genuinely feel that it is the best browser available. If you are a KDE purist, are in love with Konqueror, or simply refuse to use Firefox, you can follow the instructions in the section of this chapter entitled "Migrating Just Your Bookmarks" to export your existing favorites from Internet Explorer or bookmarks from Safari or Firefox to a file. You can then import that file into Konqueror to preserve your existing favorites/bookmark collection.**

Like Thunderbird, Firefox is sponsored by the Mozilla Foundation, so it includes the same sort of excellent settings migration tools as Thunderbird. These enable you to install Firefox on a system and use a simple menu command to import settings from an existing browser. The goal of this command is, of course, to help you switch to Firefox on any given platform, but it is also a tremendous help when migrating from one operating system to another. The next two sections explain how to migrate existing Firefox configuration information from a Windows or Mac OS X system to an Ubuntu or Kubuntu system, and how to migrate Internet Exploder settings from a Windows system to Firefox on an Ubuntu or Kubuntu system. A final section explains how to export just your bookmarks if you want to start with a clean

cache, set of cookies, and no saved passwords on an Ubuntu or Kubuntu system, but don't want to lose all of the fun and useful bookmarks that you've created in Internet Exploder or Safari.

As mentioned in previous sections, I suggest that you install Firefox and export your bookmarks on your Windows or Mac OS X system before doing backups of those systems to CD, DVD, an external drive, or network storage. This gives you a permanent snapshot of the data that you are migrating (as well as the original Web settings from your previous browser) just in case anything goes wrong and you no longer have your previous operating system to fall back on. The instructions here emphasize this by discussing things in terms of copying files from your backups, not from some other operating system. If you're dual booting Ubuntu/Kubuntu and another operating system, you can always copy files from that system.

For those who are new to Firefox, Table 12-1 shows the mappings between common Internet Explorer terms and the equivalent terms in Firefox.

TABLE 12-1

Common Internet Explorer Terms and Firefox Equivalents

Internet Explorer	Firefox
Address Bar	Location/Navigation Bar
Favorites	Bookmarks
Internet Options	Tools Options
Links Bar	Bookmarks Toolbar
Refresh	Reload
Save Target As	Save Link As
Temporary Internet Files	Cache

As a final note, remember that the version of Firefox that is available for Ubuntu/Kubuntu systems is Firefox 3, as discussed in Chapter 11. Because Firefox 3 stores bookmark and other data in a different format than previous versions of Firefox, you must make sure that the version of Firefox that you are migrating to can handle the Firefox settings data that you are migrating. If you are migrating Firefox configuration information from another system to a Ubuntu/Kubuntu system, the version of Firefox running on your Windows or Mac OS X system must be less than or equal to the version of Firefox that is installed on your Ubuntu/Kubuntu system.

Migrating from Firefox to Firefox

If you're already a Firefox user on the Windows or Mac OS X platform, congratulations on your excellent taste in Web browsers! You'll also be happy to know that it is extremely easy to migrate your existing Firefox settings (known as a profile) from any other operating system to Firefox on Ubuntu or Kubuntu systems. After installing Firefox on a Windows or Mac OS X system, dialogs that offer to import the settings from your existing Web browser(s) are displayed automatically the first time that you start Firefox.

This section assumes that you imported the settings from your previous Web browser into Firefox when you first started Firefox and have not used Internet Explorer since then. If you want to make absolutely sure that

you have gotten the latest favorites, cookies, saved passwords, and cache data from Internet Explorer, follow the instructions in the next section, "Internet Explorer to Firefox," just to be on the safe side.

Migrating from Firefox on one operating system to Firefox on another is facilitated by the fact that Firefox on any platform stores all of its bookmarks, cookies, saved passwords, and cache data in exactly the same format on all supported systems. This information, which is unique to each user, is known as your Firefox profile. The location of this profile information differs on different systems:

- **Linux and UNIX systems:** `~/.mozilla/firefox`
- **Mac OS X:** `~/Library/Application Support/Firefox`
- **Windows 2000, XP, Server 2003:** `Documents and Settings\Username\Application Data\Mozilla\Firefox`
- **Windows Vista:** `Users\Username\AppData\Roaming\Mozilla\Firefox`

This directory typically contains a Profiles directory and a profile initialization file called profiles.ini. It may also contain a directory named Crash Reports, depending upon the platform that you are using and whether Firefox has ever crashed on your system. (If the Crash Reports directory is present, you can safely delete it before doing your backups.)

Checking Your Ubuntu System First

Before migrating your existing Windows or Mac OS X Firefox profile to your new Ubuntu or Kubuntu system, make sure that you do not need to preserve anything from any Firefox profile that you may already have on your Ubuntu/Kubuntu system. If you aren't sure, try the following:

1. If you don't have a directory named `~/.mozilla/firefox` in your home directory, you don't have anything to preserve, so you can skip the rest of this list.
2. If the directory `~/.mozilla/firefox` exists in your directory, start Firefox on your Ubuntu/Kubuntu system and examine the bookmarks menu and the bookmarks toolbar. If neither of these contains any entries that you created, you don't have any bookmarks to preserve. If these do contain bookmarks that you want to preserve, follow the instructions in the sidebar in Chapter 11 entitled "Backing Up Your Bookmarks." I'll explain how to integrate your backups with the migrated Firefox settings later. Proceed to the next step.
3. Select the Edit menu's Preferences command to display the Preferences dialog, and click the Privacy icon in that dialog's toolbar. Click Show Cookies to display the Cookies dialog. If this dialog does not list any cookies, proceed to the next step. If this dialog does contain cookies, you will lose them when you migrate. Proceed to the next step.
4. Click the Privacy icon in the Preferences dialog's toolbar. If Remember saved passwords for sites is not selected, you can skip to the next step because you don't have any saved password information. If this option is selected, click Saved Passwords to display the Saved Password dialog. If this dialog is empty, close it and proceed to the next step. If this dialog does contain host and username entries, you will lose these when you migrate. Close this dialog.

TIP **If you have saved password information, you can always back this up in Firefox by clicking Show Passwords, clicking Yes on the confirmation dialog, and then either using screen capture software to capture data from this window or simply writing down any critical host/login/password information that you want to preserve.**

Finding that you have bookmarks in an existing Firefox installation that you want to preserve is easy to accomplish by backing up/exporting those bookmarks in your current version of Firefox and then restoring/importing them in the version that you migrate from another system. (See the section "Migrating Just Your Bookmarks" later in this chapter for detailed information.) However, finding cookies or saved passwords in an existing Firefox installation causes a problem. There is no easy way to merge cookie, saved password, saved form, and other cross-session data between two profiles. If you have such data in a Firefox profile on a Windows/Mac OS X system and the same type of data in an existing Firefox profile on a Ubuntu/Kubuntu system, you will essentially have to choose which Firefox profile you want to keep.

Migrating an Existing Profile

Assuming that you do not have an existing Firefox profile on your Ubuntu/Kubuntu system or don't need to preserve anything that it contains, you can do the following to migrate your existing Windows or Mac OS X Firefox profile to your new Ubuntu or Kubuntu system:

1. If you have a `~/.mozilla/firefox` directory in your home directory, rename it using a command like the following from a GNOME Terminal, Konsole, or xterm window:

   ```
   mv ~/.mozilla/firefox ~/.mozilla/firefox.OLD
   ```

 You can delete this directory later, once you're sure that you didn't somehow lose anything important.

2. Create the directory `~/.mozilla/firefox` in your home directory on your Ubuntu or Kubuntu system, using a command like the following from a GNOME Terminal, Konsole, or xterm window:

   ```
   mkdir -p ~/.mozilla/firefox
   ```

3. Copy the contents of the profile directory (identified in the location list in the previous section) from your backups to the `~/.mozilla/firefox` directory in your home directory on your Ubuntu or Kubuntu system. You can do this by dragging and dropping the contents of the profile directory from your backups to the ~/.mozilla/firefox directory using a file manager such as Nautilus, Konqueror, or Dolphin.

TIP **If you are using the Nautilus file manager on a standard GNOME-based Ubuntu desktop system, you will have to press Ctrl+H to show hidden files and directories (those whose names begin with a period) in order to navigate to the `.mozilla` directory.**

4. Start Firefox. You should see the same bookmarks, cookies, saved password, and so on, that you saw in the version of Firefox that you were running on your other operating system.

Troubleshooting Profile Migration

If you have a problem at this point, make sure that:

- The `profiles.ini` file was copied to your `~/.mozilla/firefox` directory.
- The path statement in the `profile.ini` file identifies a directory that is also located in your `~/.mozilla/firefox` directory. This is usually a directory with a name like *`xxxxxxxx`*`.default`, or a directory called Profile that contains a directory with a name like *`xxxxxxxx`*`.default`.

- The `~/.mozilla/firefox` directory and all of its subdirectories are owned and are readable by you. You can ensure this by executing commands such as the following from a GNOME Terminal, Konsole, or xterm window:

```
sudo chown -Rv username ~/.mozilla/firefox
sudo chmod -R u+r ~/.mozilla/firefox
```

If either of the items in the first two bullets is missing, re-copy the contents of your Windows or Mac OS X profile directory to your `~/.mozilla/firefox directory`. Make sure that you copy the contents of your Windows or Mac OS X profile directory, and not the directory itself. Your `~/.mozilla/firefox` directory must contain a `profiles.ini` file that identifies a valid directory where your personal profile data is located.

Migrating from Internet Explorer/Safari to Firefox

If you have not used Firefox on your Windows or Mac OS X system, you can take advantage of its configuration import and conversion capabilities to simplify migrating your Internet Explorer or Safari settings to Ubuntu or Kubuntu. Firefox is a useful migration tool because the first time that you start it on a new system, it displays a series of dialogs that offer to import the settings from your existing Web browser(s). These settings are automatically converted into the formats used by Firefox. You can then follow the instructions given in the previous section for migrating Firefox settings from one system to another.

NOTE **If you are migrating bookmarks from Internet Explorer or Safari to Konqueror, you will want to follow the instructions in the section "Migrating Just Your Bookmarks" rather than the instructions in this section. This section applies only if the browser to which you want to migrate your current browser settings is the Firefox browser. You can install and use Firefox on Kubuntu systems (which I strongly suggest), but if only Konqueror will do for you, skip to the section "Migrating Just Your Bookmarks."**

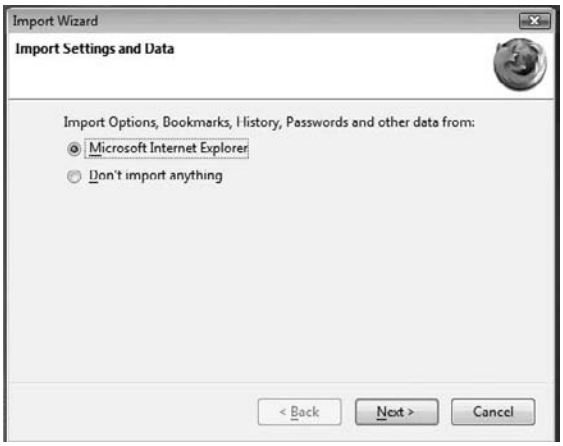
The first step in this process is, of course, downloading and installing Firefox on your Windows or Mac OS X system. You can get the latest official version of Firefox for Windows or Mac OS X from `www.mozilla.com/en-US/firefox/all.html`. Firefox for Windows downloads as a single executable, while Firefox for Mac OS X downloads as a disk image (.dmg) file that you must mount in order to find the installer.

NOTE **This section uses screen captures from a Windows system. The dialogs are essentially the same on a Mac OS X system. Any significant differences between dialogs and options are described in the text.**

Once you have downloaded Firefox, install it using your system's standard installation mechanism. The summary dialog at the end of the installation process provides a checkbox that enables you to start Firefox immediately. Click Finish to do so. If you are using a Windows systems, a dialog like the one shown in Figure 12-12 displays. If you are using a Mac OS X system, a similar dialog with different options displays.

FIGURE 12-12

Selecting the browser to import settings from



If you are using a Windows system, select the radio button beside Microsoft Internet Explorer and click Next. If you are using a Mac OS X system, select the radio button beside Safari and click Next. A dialog like the one shown in Figure 12-13 displays.

FIGURE 12-13

Specifying your Firefox home page



If you are happy with the home page that you've been using in Internet Explorer, click the radio button beside Import your home page from Internet Explorer. If you would prefer to use a Firefox home page, make sure that the Firefox Start . . . radio button is selected. Click Next to proceed. The import process begins. When it completes, a dialog like the one shown in Figure 12-14 displays.

FIGURE 12-14

The final import status dialog in Firefox



This dialog identifies all of the items that have been imported from Internet Explorer or Safari into Firefox. Click Finish to close this dialog and start Firefox. Favorites that have been imported from Internet Explorer will be listed on a menu called From Internet Explorer, and bookmarks that have been imported from Safari will be located on a menu called From Safari. Whichever of these menus you see will be located on Firefox's standard Bookmarks menu.

Once you have imported your browser settings from Internet Explorer or Safari into Firefox on a Windows or Mac OS X system, the remainder of the process of migrating that configuration data to Firefox on your new Ubuntu/Kubuntu system is the same as that explained in the previous section of this chapter, "Migrating from Firefox to Firefox."

Migrating Just Your Bookmarks

The previous sections discussed ways of migrating all of your Web settings from one computer system to another, primarily from Internet Explorer or Safari to Firefox on a Linux system. This can be done pretty easily using Firefox as a migration tool, as explained in the previous section. By default, migrating via Firefox preserves everything — favorites/bookmarks, cookies, saved passwords, and so on. However, you may not always want to preserve everything. Using a Web browser for a long time is like living in the same house for a long time — you tend to accumulate stuff and you may not even be aware of everything that you've accumulated. In a Web browser, this can mean that you have saved passwords to sites that you no longer visit, cached data that was useless long ago, collected cookies for sites that you last visited to purchase Y2K survival supplies, and so on. While you probably don't care about much of this stuff (and therefore won't miss it if you don't migrate it), one thing that you do care about and probably use is your collection of links to favorite sites, known as "favorites" to Internet Explorer users and as "bookmarks" to the rest of the known universe.

This section explains how to save your Internet Explorer or Safari bookmarks to an HTML file that you can then import into the Firefox and Konqueror browsers. As suggested in the previous paragraph, you may want to migrate only your bookmarks when moving to a new system as a sort of mandatory

housekeeping exercise. However, being able to create HTML files containing your bookmarks can also be very useful if you need to use other computers and operating systems but want to have your familiar bookmarks there. Some businesses mandate using a specific operating system, and some Web sites require specific browsers (sadly enough). In these cases, being able to export bookmarks from one browser and import them into another provides a usable, but somewhat crude, way of seeing the same bookmarks in multiple browsers.

A popular solution to the problem of seeing the same bookmarks from multiple browsers is to make your bookmarks available on the Web. You can do this by making an HTML file containing your bookmarks available on your Web site (as I do), or by using various tools that enable you to store your bookmarks on specially designed Web sites. For a huge list of Web-based bookmark management tools, see `www.dmoz.org/Computers/Internet/On_the_Web/Web_Applications/Bookmark_Managers`.

Regardless of why you want to do it, being able to save your current favorites/bookmarks to an HTML file is very convenient. Whether you're doing it for migration purposes, quick and dirty backups, personal convenience, or to share your favorite Web time sinks with your friend, the next few sections explain how to export from and import to different popular browsers.

TIP **If you are using a browser other than the ones discussed in this section, see the MozillaZine Web page on importing bookmarks, `http://kb.mozillazine.org/Import_bookmarks`. This page also discusses how to export bookmarks from a number of browsers other than IE or Safari.**

Exporting Internet Explorer Bookmarks

Both Internet Explorer 6 and 7 provide a convenient Import and Export dialog (shown in Figure 12-15) that makes it easy to export your favorites to a file that you can then import into other browsers on other systems — only the location of the command used to display this dialog has been changed to confuse the user.

FIGURE 12-15

Internet Explorer's Import and Export Wizard



To display this dialog in IE6, select the File menu's Import and Export menu item. To display this dialog in IE7, click the Add to Favorites icon (second from left in your default toolbar), and select the Import and Export menu command. Click Next to proceed, which displays the dialog shown in Figure 12-16.

The dialog shown in Figure 12-16 enables you to select the type of information that you want to export from Internet Explorer: Favorites, Cookies, or RSS Feeds (the latter is only shown if you are exporting from IE7). To export your Favorites as an HTML file, select Export Favorites and click Next to proceed. The dialog shown in Figure 12-17 displays.

FIGURE 12-16

Selecting what to export from Internet Explorer

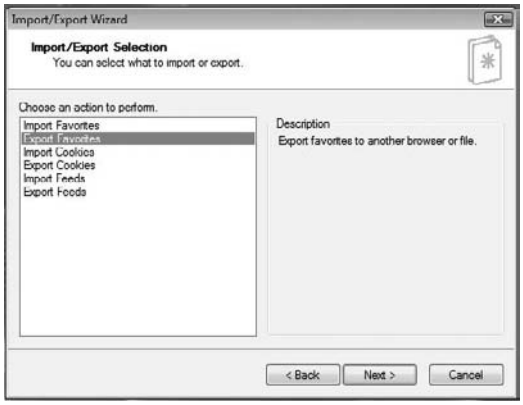


FIGURE 12-17

Selecting the Favorites folder(s) to export items from Internet Explorer



The dialog shown in Figure 12-17 enables you to select the Favorites folders that you want to export as HTML. To export your entire collection of Favorites, select the top-level Favorites folder, as shown in Figure 12-17. Click Next to proceed. The dialog shown in Figure 12-18 displays.

FIGURE 12-18

Selecting where to save exported items from Internet Explorer



The dialog shown in Figure 12-18 enables you to specify the location in which the file containing your favorites (ironically called `bookmark.htm`) will be saved, or to optionally export them directly to another application. Because our goal is to save them to a file (which is the default action), leave the Export to a File or Address radio button selected and modify the name of the file if you want to use a name other than the default. Click Next to proceed.

A final dialog displays, which requires that you click Finish to begin the export process. A confirmation dialog displays when the export process is complete. Click OK to close the confirmation dialog, which will also close the Import and Export dialog.

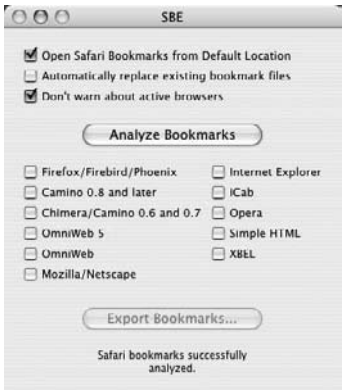
Copy the file containing your favorites to the system where the browser to which you want to import them is running. If this system does not exist yet or is currently inaccessible, make sure that the file that you just created is saved to a location that will be included in the backups that you will be creating of your Windows system.

Exporting Safari Bookmarks

Strangely enough, the Safari browser that Apple provides for Mac OS X systems does not automatically provide the ability to save its bookmarks to an HTML. The good news is that Simone Manganelli has written an excellent (and free) application called the Safari Bookmark Exporter (`http://homepage.mac.com/simx/sbe.html`) that does just that. After downloading the disk image file for this application, you can either install it or run it directly from the image. A dialog like the one shown in Figure 12-19 (captured on a Mac OS X system) displays.

FIGURE 12-19

The initial Safari Bookmark Exporter dialog



If Safari is currently running on your OS X system, click the Don't warn about active browsers button. Next, select an output target by clicking one of the browser/file format names below the Analyze Bookmarks button — your selection primarily determines where the Safari Bookmark Exporter will try to save the file containing your bookmarks. I typically select Firefox/Firebird/Phoenix just in case this option does anything special for Firefox, although Simple HTML is also a reasonable choice (especially if you are planning to import your bookmarks into Konqueror). Click Export Bookmarks to display a Finder dialog that enables you to identify where you want to save your bookmarks. If you selected Firefox/Firebird/Phoenix as your browser, the suggested location will be to save the output file to your Firefox profile directory, which you may not want to do because it will overwrite any file of the same name. I typically save these files to my home directory with a name like `safari-bookmarks-export.html`, which makes it pretty clear what the file contains.

Quick Hack for Saving Safari Bookmarks

If you don't want to download and install the Safari Bookmark Exporter application for a one-time effort, you can use some internal Safari hacks to save your Safari bookmarks to an HTML file. Because this is indeed a hack, it may not work in all newer versions of Safari, but it's worth a try:

1. Exit Safari and start the Terminal application.
2. Execute the following command in the Terminal window:

```
defaults write com.apple.Safari IncludeDebugMenu 1
```

3. Start Safari.
4. Select the new Debug menu's Export Bookmarks command and identify the location to which you want to save the Safari `Bookmarks.html` file.

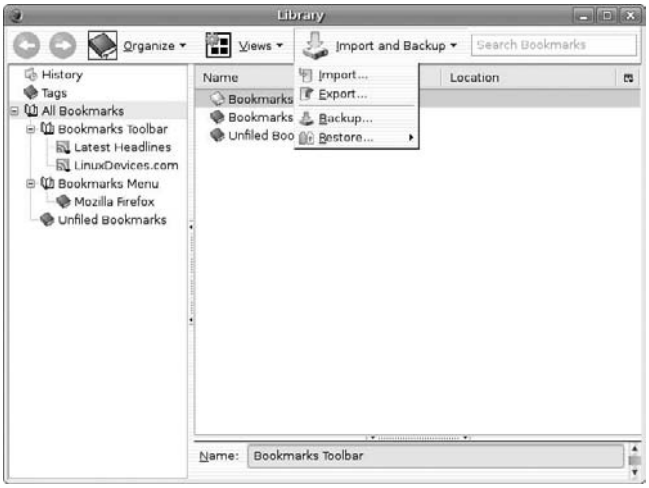
You can then import this file into the Firefox or Konqueror browsers, as described later in the chapter.

Importing Bookmarks into Firefox

To import an HTML file containing bookmarks into Firefox, first select the Bookmarks menu's Organize Bookmarks command to display the Library dialog. Once this dialog displays, click the Import and Backup menu to display the list of available commands, as shown in Figure 12-20.

FIGURE 12-20

Commands on the Library dialog's Import and Backup menu



Select Import from this menu to display the Firefox Import Wizard, as shown in Figure 12-21.

FIGURE 12-21

The Firefox Import Wizard



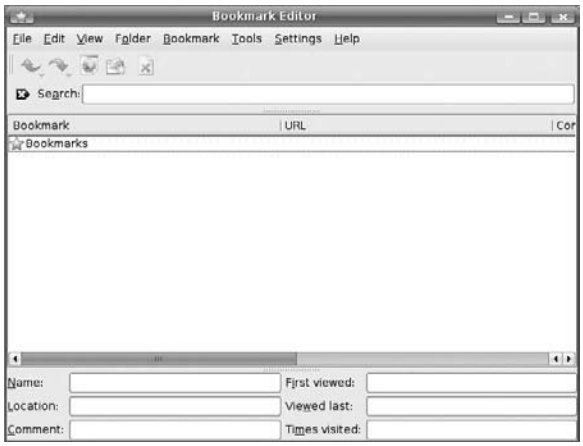
Click Next to proceed. A standard file navigation dialog displays, in which you can browse to and select the file containing the bookmarks that you want to import. Click Open to open that file. All of the bookmarks contained in that file will be added to the Bookmarks menu folder in Firefox. Click the Close box to dismiss the Library dialog and return to Firefox.

Importing Bookmarks into Konqueror

To import an HTML file containing bookmarks into Konqueror, first select the Bookmarks menu's Edit Bookmarks option to display the Bookmark Editor dialog, as shown in Figure 12-22.

FIGURE 12-22

Konqueror's Bookmark Editor



TIP **You can also display this dialog without going through Konqueror by executing the `kedit-bookmarks` command from a KDE Konsole or xterm window.**

Select the File menu's Import command to display the submenu shown in Figure 12-23, and select the appropriate command from this menu.

The primary differences between these commands are the default location in which they expect to find bookmark files and the format of these files. I always select the Import Mozilla Bookmarks option because it enables you to select an HTML file to import and has always worked for me when importing bookmark files from different versions of Internet Explorer and Safari, which are the most common migration scenarios that I've been involved with.

After selecting a specific import type, a navigation dialog displays, in which you can browse to different directories and select the HTML file that you want to import. After selecting an HTML file containing bookmarks, click Open to open that file. A dialog like the one shown in Figure 12-24 displays.

FIGURE 12-23

Supported bookmark import options in Konqueror

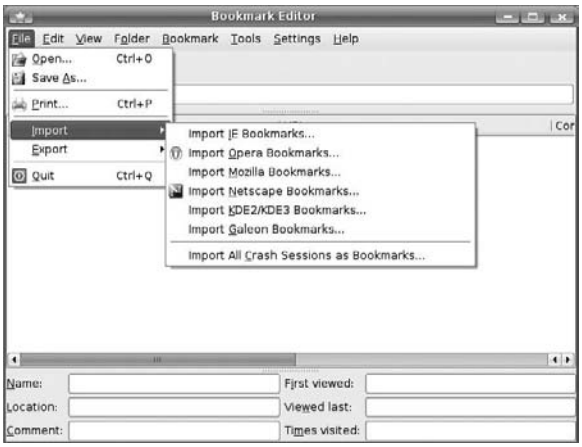


FIGURE 12-24

Storing imported bookmarks in Konqueror



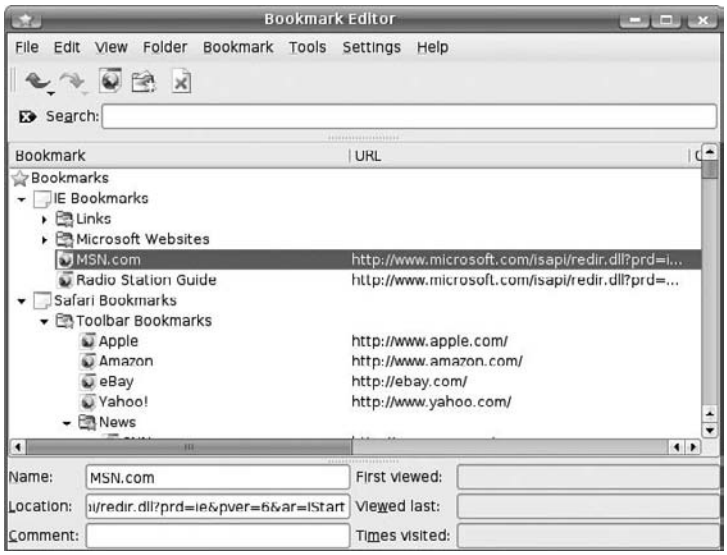
This dialog enables you to specify whether the bookmarks that you are importing will replace your current Bookmarks folder or whether they should be created as a new subfolder. I always select As New

Folder, first because I don't want to lose my existing bookmarks, and second because this makes it easy to identify imported bookmarks, which is especially useful when you are manually reorganizing your bookmark collection.

After clicking As New Folder, your bookmarks are imported and put into a new subfolder whose name is based on the menu command that you selected for import. Because I always use the Import Mozilla Bookmarks menu option, I typically rename the new subfolder to identify the true origin of the bookmarks by right-clicking on the subfolder and selecting Rename from the context-sensitive menu. Figure 12-25 shows the Bookmark Editor displaying my original bookmarks folder, and sets of bookmarks that I have imported from both Internet Explorer (the IE Bookmarks section) and Safari (the Safari Bookmarks section).

FIGURE 12-25

Imported bookmarks shown in the Bookmarks Editor



Working with Windows Application Data

The previous two sections of this chapter discussed how to migrate specific types of data — your e-mail and browser-related data — from another operating system to your new Ubuntu/Kubuntu system. Although these are probably extremely important items for you, they're only the tip of the iceberg in terms of all of the data that you probably have on your previous computer system.

The general-purpose section entitled "Backing Up Your Personal Data" stressed the need for backing up all of your personal information, regardless of where it is located, from any computer that you are no longer planning to use — especially if you are planning on erasing or overwriting the disk on which your

current data is located. In general, I suggest that you back up your data and preserve your existing disk whenever possible, until you are sure that you haven't accidentally left some critical files or directories out of your backups. This is easy to do in desktop systems because you can pull your old disk and replace it with a new one on which you can install Ubuntu/Kubuntu, but if that's not possible for one reason or another in your case, please verify your backups before proceeding.

Copying files and directories from CDs, DVDs, external drives, or network storage locations that contain your backups is a simple command-line or drag-and-drop operation once your Ubuntu/Kubuntu system has mounted the CD, DVD, or networked storage device. CDs and DVDs should automount when inserted into your Ubuntu/Kubuntu system unless you've disabled automounting, but you will probably have to use the command-line or a graphical utility such as Ubuntu's Places ➪ Connect to Server dialog or Kubuntu's Konqueror ➪ remote:/ screen to mount networked backups so that you can access them via Samba or NFS.

Copying files and directories is one thing — being able to actually work with them is another. Although I am a lifelong UNIX fanatic and have been a Ubuntu devotee since the first time I reviewed it years ago, here are the harsh truths about using any Linux system to work with files that have been created by Windows or Mac OS X applications. Depending on the type of file and the version of the application that created it, you may:

- Not be able to display, print, or edit your existing files in any way
- Not be able to do as good a job of displaying, printing, or editing your existing files
- Be able to do everything you could before, but using free Linux or Web-based applications
- Need to install your Windows software on your Linux system and run it using emulation, virtualization, or a similar solution

The first two of these are horror stories that can happen. The applications developed by most software vendors traditionally use their own data file formats. If that software vendor does not support Linux, whether or not Linux software exists that can read and write the data files produced by that vendor's applications depends on whether the data file format use by that application is well-known, open, or has been reverse engineered. It also depends on whether the Linux developers of similar applications care enough to try to support reading and writing files in that format.

As much as it pains me to say it, if you are doing critical work in a Windows or Mac OS X application, you should use the Ubuntu Live CD or a scratch Ubuntu system to test if or how well you can work with your existing data files on a Linux system. Most of the chapters in the first part of this book are organized by common types of tasks and the Linux applications that you can use to accomplish them — writing documents, working with spreadsheets, creating presentations, and so on.

The following are a few commonly encountered types of files beyond those discussed in earlier chapters and the Linux applications that you can use to work with them:

- **AVI, MPG, SWF:** My favorite Linux player for all of these is the mplayer application, but applications such as Totem on Ubuntu and Kaffeine on Kubuntu systems are popular and improving with each release.
- **MP3 files:** You can use the EasyTAG application to edit the ID tags for MP3 files. The mpg123 and mpg321 applications are command-line utilities for playing MP3 files, which you can also play in graphical applications such as Rhythmbox and Helix Player under GNOME, Amarok and Kaffeine under KDE, and many more. You can encode various formats to MP3 using the

graphical Sound Converter or command-line lame applications on Ubuntu, or the graphical AudioKonverter application on Kubuntu.

- **DOC, XLS, PPT:** You can use the OpenOffice.org tools to open the files that are created by Microsoft Office applications. Microsoft's new XML-oriented DOCX, XSLX, and related formats are only supported by OpenOffice.org version 3.0 (and greater) tools.
- **PDF files:** Ubuntu provides the Evince application to display PDF files, and Kubuntu provides the kpdf application under KDE 3 and the okular application if you are running KDE 4. All of these applications are read-only applications. Most Ubuntu and Kubuntu applications enable you to create PDF files by selecting the Print to File option and specifying PDF output. The pdfedit application, which can be found in the Ubuntu repositories, can be used to edit existing PDF files.

As you can see, this list could go on forever. For a great summary of the Linux applications that are equivalent to popular Windows applications, see sites such as the following:

- Linux Alternatives Project: `www.linuxalt.com/`
- Linux Application Finder: `www.linuxappfinder.com/alternatives`
- Linux Equivalents to Windows Software: `www.linuxlinks.com/article/20070701111340544/Equivalents.html`
- Linux Software Equivalent to Windows Software: `http://wiki.linuxquestions.org/wiki/Linux_software_equivalent_to_Windows_software`
- Table of Equivalents/Replacements/Analogs of Windows software in Linux: `www.linuxrsp.ru/win-lin-soft/table-eng.html`

In addition to these Linux-specific sites, the Open Source Alternatives site (`www.osalt.com/`) lists equivalent applications for Windows software that run on other open source operating systems (including Linux, of course).

For information about using emulation, virtualization, and similar techniques to enable you to run actual Windows software on your Ubuntu/Kubuntu system, see Bonus Chapter 7. Yes, it can be done!

Summary

The wealth of software available for Linux makes it quite conceivable to make Ubuntu or Kubuntu your "daily driver" operating system — the one where you get all of your work done. But what about all of the mail, browser favorites/bookmarks, and data files that you've accumulated on your existing Windows or Mac OS X systems? This chapter explained different ways of migrating e-mail and browser settings from a Microsoft Windows or Mac OS X system to an Ubuntu or Kubuntu system, discussed good techniques for backing up your Windows or Mac OS X system to facilitate migration, identified some tasks that you should do before doing your final backups, and discussed critical directories where applications such as Internet Explorer, Outlook, and Outlook Express store their data. It also discussed other considerations for complete migration to a Ubuntu/Kubuntu system, such as making sure that you can find applications that will enable you to work with your existing data files.

Chapter 13

Sending and Receiving Instant Messages

E-mail has revolutionized how people communicate today. E-mail is the carrier of most business communications, replacing paper memos, faxes, and most information-sharing meetings, and enabling geographically disparate offices and personnel to keep in touch and know what each is doing. E-mail has revolutionized personal communications, too. Thanks to e-mail, I've been able to maintain friendships that would otherwise have atrophied because I'm not the best letter writer or snail mail user. I can send e-mail at any time, from any computer, and know that it will be waiting for the recipient after a very short time.

"Waiting for the recipient" identifies a key shortcoming of e-mail. E-mail was designed as an asynchronous communications mechanism, and does a great job at that. I can send you e-mail and you can respond at your convenience. Unfortunately, "at your convenience" isn't always good enough. If I have an extra ticket to a baseball game that starts in a few hours, I need to know whether you want it now. If you send an instant message to someone who is running an instant messaging client, they are usually alerted that a new message has come in—you don't have to wait for them to get around to checking their mail. A related issue is that e-mail isn't really conversation-oriented. If I send e-mail with several questions, you can respond to my e-mail at the top or bottom, or intersperse your responses, but a back-and-forth exchange of conversational messages quickly gets hard to read, at best.

Instant messaging, where you can open a direct connection between one or more people and either send a single message or have an online conversation, has been around for quite a while, but is still increasing in popularity as both a personal and business communication mechanism. This chapter provides a little background on instant messaging, explains what you need to do in order to participate in instant messaging, discusses the most popular instant messaging clients on Ubuntu and Kubuntu systems (Pidgin and Kopete, respectively), discusses Internet Relay Chat, and provides some tips on resolving instant messaging problems.

IN THIS CHAPTER

- Instant messaging history
- Getting an IM account
- Using Pidgin on Ubuntu
- Using Kopete on Kubuntu
- Using IRC
- Other IM clients
- Troubleshooting

Instant Messaging Overview

Believe it or not, people were actually online even before the Internet. Bulletin-board systems (commonly referred to as BBSs) were a popular early personal computing phenomenon, consisting of computers running specialized, multi-user software that users could connect to using a modem, and on which users could read and send both public and private messages using a terminal-based interface. (Many BBS systems also supported inter-BBS e-mail by forwarding messages from one system to another using modems and forwarding networks such as FidoNet.) Commercial dialup/BBS services such as America Online (AOL), CompuServe, Prodigy, and many others provided centralized locations on which people could share resources and, more important, instantly communicate with each other in shared areas, commonly referred to as "chat rooms," conferences, or echoes (FidoNet), or online forums. As the availability, popularity, and amount of content on the Internet increased in the early 1990s, the reason for these centralized services became less important—after all, once you can contact anyone anywhere via e-mail and have access to the entire Internet, why bother limiting yourself to a single online community? The entire planet is your potential community.

As discussed in the introduction to this chapter, e-mail is cool but is inherently message-oriented rather than conversation-oriented. UNIX and UNIX-like systems have always had real-time text communication software, starting with the `write` command in the original research edition of UNIX, which enabled you to immediately send a one-way message to another user on the terminal or other device that they were using. (Enhanced versions of the `write` command, such as Orville Write—`http://unixpapa.com/write.html`—supported multi-user and multi-way communication, but were not part of the standard UNIX distribution.) Early UNIX distributions also provided the `talk` command, which enabled you to have real-time text conversations between two users by splitting your terminal window into two sections in which you could exchange messages back and forth. (I'm told that the first version intermingled text as both users typed, but I'm not quite that old.) I remember using the `talk` command in the early 1990s to help an acquaintance in Australia debug some problems he was having with one of his machines, while I sat in my office in Pittsburgh. Good times…

Outside the UNIX-only universe, the increased availability of the Internet led to the development of many cross-platform and admittedly friendlier instant communication mechanisms. The idea of a shared, instant communication area is always attractive for its immediacy. Modern Internet-based instant communication was born in 1988 with the creation of IRC (Internet Relay Chat), written by Jarkko Oikarinen (`www.ircnet.org/History/jarkko.html`). IRC, still in wide use today, enables instant communication between users by providing IRC channels (much like chat rooms) that users connect to, at which point their messages are seen by anyone on the channel. IRC is a client/server system, where users connect to an IRC server in order to join an IRC channel that is supported by that server. While IRC enables user-to-user communications, it is primarily a user-to-channel message system reminiscent of chat rooms on bulletin-board systems. IRC was born on UNIX systems (big surprise), but clients quickly appeared for the Macintosh, Microsoft Windows, and DOS environments.

NOTE **Regardless of how well it fits the IM model, IRC is still popular, quite fun, and occasionally even useful. For example, the #ubuntu IRC channel is a great source of immediate help for Ubuntu (although there's also a lot of noise there). If you're using the standard GNOME-based Ubuntu desktop, you can communicate via IRC using Pidgin. If you're using Kubuntu, you can communicate via IRC using Kopete. Because IRC is somewhat different than instant messaging per se, it's discussed in a separate section later in this chapter. See the section of this chapter entitled "Using IRC" for general information about IRC and specific information about using IRC on Ubuntu and Kubuntu systems.**

What most people think of as instant messaging today first appeared in 1996, when Mirabilis introduced its ICQ ("I Seek You") instant messaging software, which is still popular today (although now owned by Time Warner/AOL). ICQ is a client/server instant messaging system, where desktop users use ICQ client software to instantly exchange messages with other users through centralized ICQ servers. This is the same model used by most instant messaging software today.

ICQ was amazing when it first appeared and is still quite cool, although certain aspects of it show its age. All ICQ users are primarily identified by the User Identification Number (UIN), which is a bit kludgy nowadays. When's the last time you thought to yourself "Wow, I've got to tell 481241598 about this!"? Although newer versions of ICQ allow identification via your e-mail address, most modern instant messaging systems let you create your own user name, which must be unique within that system but can at least be reminiscent of your name or how you want yourself to be perceived online. Mine are usually something reminiscent of my name, but identities such as "topdog" and "longdistancelover" are more attractive to others. Like any good technology, it's all up to you.

Most modern instant messaging applications provide core features such as the following:

- **Instant messages:** Exchange text messages with a friend who is online via a computer or cell phone
- **Chat:** Have an electronic conversation with multiple people at the same time
- **Exchange files:** Send audio, video, and other files to one or more contacts

Instant messaging is lots of fun, and can even be useful in business communications. Many companies today rely on instant messaging for situations when instant contact is important or simply to save you from having to keep walking over to someone else's desk and physically interrupting them (which can be tricky if you're in different parts of the country or the world). As a general rule, people seem to be less verbose when sending instant messages, which is often a good thing if you just need a bit of info to get on with your work.



One important thing to remember when using instant messaging, especially for business purposes, is that the protocols used by many instant messaging services are not guaranteed to be secure. You may not want to send someone important business information, your credit card numbers, or your passwords via an IM client.

Nowadays, you'll probably find some occasion to IM someone at some point. (As you can see, instant messaging is so popular that it's both a noun and a verb!) If you're not already registered with a system that supports instant messaging, see the next IM section for information about registering with one of the many free services that are available. If you already have an instant messaging account and just want to know how to join the IM scene from your desktop or laptop, skip ahead to the section of this chapter that discusses the standard instant messaging client for the Ubuntu distribution that you are using: either Pidgin on Ubuntu desktop systems, or Kopete on Kubuntu systems. Of course, many other IM clients are available for Ubuntu and Kubuntu systems—the section entitled "Other Instant Messaging Packages," later in this chapter, provides an overview of some other popular IM clients for GNOME and KDE systems.



Fans of doing everything inside Emacs may want to see the article at `www.instantmessagingplanet.com/public/article.php/3742411`, which describes how to do instant messaging (and twittering) inside Emacs. Although I love Emacs like a brother, I'm not quite that dedicated, but your mileage may vary.

Getting an Instant Messaging Account

To exchange instant messages with someone else, you need to be registered with some centralized system that supports instant messaging. You can get an instant messaging account by doing any of the following:

- Registering with an instant messaging-focused service such as AIM, ICQ, or Jabber. (AIM will also toss in a free e-mail account, but their focus is still instant messaging. E-mail is just gravy.)
- Registering with the instant messaging services provided by companies such as AOL, Google, Microsoft, Yahoo, and many others.
- Using a free Internet mail service such as AOL, Yahoo, Google, Microsoft Hotmail/Live, and so on, which also provide IM access.

One thing to consider when selecting or signing up for an IM service is what other IM services it can connect to, if any. You'd think that instant messaging providers such as AOL, Yahoo, Microsoft, and others would realize that it was useful for them to be able to support receiving instant messages from each other, but this is woefully not the case.

Most of the free instant messaging services focus on receiving messages from other members of the same service—compatibility between them is limited because all of these services use their own protocols to communicate between your desktop client and the instant messaging servers that are run by these companies. This isn't an issue for Linux users because Ubuntu and Kubuntu IM clients such as Pidgin and Kopete are what is known as multi-protocol clients, meaning that they understand the protocols used to communicate with different IM services, and thus enable you to exchange instant messages with almost anybody. Windows and Mac users aren't so lucky—although multi-protocol instant messaging clients are available for these platforms, instant messaging clients such as MSN Messenger, Yahoo Messenger, Apple's iChat (AIM and Jabber), and others primarily focus on communicating with other users of those IM services, although some offer limited interoperability. They all support some level of SMS (Short Message Service) compatibility so that they can send instant messages to cell phones.

Multi-Protocol IM Clients for Other Systems

You should encourage your Macintosh and Windows friends to install and use multi-protocol IM clients on their systems, so that they can exchange instant messages with anyone regardless of the IM service that they use.

- **Windows:** A version of the Pidgin IM client provided on Ubuntu systems is available for Windows (`www.pidgin.im/download`), and is therefore the package that I'd recommend. Another popular open source IM client for Windows is Miranda-IM (`www.miranda-im.org`). As far as proprietary, closed-source IM clients for Windows go, the most popular native multi-protocol IM client for Windows systems is probably Trillian (`www.ceruleanstudios.com`), which is also being ported to Mac OS X. Another popular multi-protocol client for Windows is Qnext (`www.qnext.com/universal_messenger.shtml`).
- **Mac OS X:** Alas, Pidgin is not available for Mac OS X, so Mac OS X users should get Adium X (`www.adiumx.com`), the most popular multi-protocol client on that platform.

For detailed and overwhelming information about IM clients, see the comparison of instant messaging clients on Wikipedia at `http://en.wikipedia.org/wiki/Comparison_of_instant_messaging_clients`.

If you're new to instant messaging, the big question is "Whom should I sign up with?" The best answer to this question is "Where are your friends?" Although multi-protocol clients such as Pidgin and Kopete enable you to send and receive instant messages from multiple services, you still have to be registered with those services for people to be able to send you instant messages through them. Microsoft and Yahoo enable their IM users to send instant messages to each other, but interoperability is still a relatively uncommon notion for instant messaging servers.

Although using a limited-protocol, closed-protocol site such as AIM, MSN Messenger, or Yahoo only encourages the persistence of proprietary protocols and associated tools, being able to talk to your friends is probably more important than philosophy. The good news is that because most instant messaging accounts are free, you can sign up with as many instant messaging services as you want and then use a multi-protocol client to communicate with all of your friends at any of them.

NOTE **If you're already using a free Internet mail service like AOL, Microsoft Live/Hotmail, Google, or Yahoo, you probably already have an instant messaging account through the same service. Log in to your webmail account and look for an FAQ (Frequently Asked Questions file) or general information about IM support. In most cases, your IM identify is the same as your e-mail address, without any domain information. For example, if your e-mail address is wvonhagen@yahoo.com, your Yahoo IM identity is almost certainly wvonhagen.**

TIP **If you are using Kubuntu, you can register for any of the services that are supported by KDE's Kopete instant messaging application through Kopete itself. See the section later in this chapter entitled "Instant Messaging on Kubuntu with Kopete" for more information.**

The following is a list of some popular locations where you can sign up for free instant messaging accounts:

- **AIM:** AOL's acquisition of Mirabilis's ICQ software made AIM (AOL Instant Messenger) the original leader in the instant messaging space, and it's still one of the most popular IM sites around. You can get a free AIM identity (known as a screen name because of its roots in AOL) from `www.aim.com`, which comes with an equivalent aim.com e-mail account. Similarly, you can get an AIM account by registering for free e-mail with AOL (`www.aol.com`). The underlying communications protocol used by AIM is called Oscar.
- **Google:** Far be it from Google to miss out on any aspect of online life. Google Talk (`www.google.com/talk`) is Google's instant messaging platform, which you can access through the same Google account that you use for Gmail. If you do not have a Google account, launching Google Talk provides links through which you can create one. Not too surprisingly, Google Talk uses a browser interface for instant messaging and chats, so it works almost anywhere. Google Talk uses XMPP (Extensible Messaging and Presence Protocol) as its underlying protocol.
- **ICQ:** Although still focused on numeric IDs, ICQ popularized instant messaging, and is still a popular site. You can register for an ICQ account at `https://www.icq.com/register`. Because AOL bought ICQ long ago, ICQ also uses Oscar as its underlying protocol.
- **Jabber:** Jabber is an open instant messaging platform where, as its PR material states, "…your messages are not monitored by a big company like AOL, Microsoft, or Yahoo." Jabber uses the open XMPP (Extensible Messaging and Presence Protocol), which is also used by Google Talk. To register with Jabber, your IM client will automatically offer to create an account for you when you use any of the sites that use Jabber (see `www.jabber.org/im-services` for a list).
- **Microsoft:** Microsoft provides free Microsoft Live e-mail accounts (formerly Hotmail), which come with a free Windows Live Messenger (formerly known as MSN Messenger) account because all of your Microsoft accounts share the same identity. See `http://login.live.com` to sign

up for an account. I know that you'll be surprised to learn that Microsoft uses its own, proprietary IM protocol for instant messaging.

- **Yahoo:** Yahoo has provided an instant message service for quite a while, and its Yahoo Messenger software is quite popular. You get instant messaging capabilities when you sign up for a Yahoo e-mail account (`www.yahoo.com`). Like Microsoft, Yahoo uses its own protocol for instant messaging.

These are just the big names in the IM space. See the subsequent sections of this chapter on Pidgin and Kopete for information about the protocols that are supported by each of these instant messaging clients.

Although I'll admit to being a little afraid of Googlezilla, if I were to suggest one of these sites for your own IM pleasure, I'd vote for Google Talk, simply because you can then use Gmail and the many other useful Google tools. Because it is open source software that uses a standard protocol, Jabber is a close second, especially if most of your IM friends are also Linux fans. If not, Yahoo and AIM are your best bets.

Instant Messaging on Ubuntu with Pidgin

Pidgin (`www.pidgin.im`) is the new name of a popular instant messaging application for Linux that was originally known as GAIM (GTK+ AOL Instant Messenger). This name was changed to Gaim due to pressure from AOL, which still wasn't good enough once AOL trademarked the term "AIM," so there you go—Pidgin it is. The name is doubly clever, referring both to a simplified language used to communicate between groups who do not have an entire language in common (i.e., "pidgin english") and also referring to the "pigeon," a bird that historically was used to send messages.

Pidgin is the most popular multi-protocol instant messaging client on GNOME-based systems, and is installed as part of a default Ubuntu desktop system. Pidgin supports a huge number of instant messaging protocols, enables you to set up chat rooms for use by multiple users, and even supports file transfers.

Some terms that you'll encounter when using Pidgin are Finch, which is a terminal-oriented version of Pidgin, and Purple, which is the library (actually libpurple) that enables Pidgin and Finch to connect to multiple instant messaging systems, manages your contacts accounts, and so on. Purple is also used by other instant messaging clients, such as Adium X on Mac OS X, ScatterChat, Meebo, and so on.

To start Pidgin on an Ubuntu desktop system, select the Applications ➪ Internet ➪ Pigeon Internet Messenger menu item. An icon for the Pidgin application displays in your panel's system tray/notification area, and either Pidgin's Buddy List window or an account configuration dialog displays on your screen. Get ready to IM!

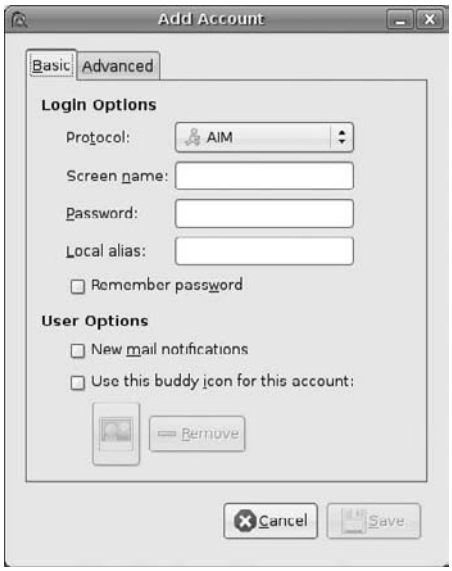
Configuring Instant Messaging Accounts in Pidgin

The first time that you start Pidgin, it displays a dialog letting you know that you haven't configured any instant messaging accounts yet, and offering to let you configure your first one. Click Add to proceed. The Add Account dialog displays, as shown in Figure 13-1.

NOTE **If you have already defined an IM account in Pidgin and want to add another one, you can do so at any time from Pidgin's Accounts dialog, which you can display by selecting the Accounts ➪ Manage menu option from the Pidgin Buddy List window. Click Add to define a new Pidgin account from this dialog. This displays the same dialog shown in Figure 13-1.**

FIGURE 13-1

Configuring an IM account in Pidgin



Click the Protocol drop-down menu to select the instant messaging protocol that you want to use, which essentially means the instant messaging service where you have an instant messaging account.

The Pidgin instant messaging client provided on Ubuntu Linux systems supports the following instant messaging protocols and associated sites:

- **AIM:** (See the previous section of this chapter for information about creating an AIM account.)
- **Bonjour:** The new name for Apple's Rendezvous, the automatic configuration and location protocol used by iChat. Rendezvous is an example of zero configuration (zeroconf) software, where clients use special packets and network information to locate each other.
- **Gadu-Gadu:** A Polish instant messaging client and associated protocol. Gadu-Gadu uses numeric identifiers as user identities, much like ICQ. You can create a Gadu-Gadu account from within Pidgin (`www.gadu-gadu.pl`).
- **Google Talk:** (See the previous section of this chapter for information about creating a Google Talk account.)
- **Groupwise:** Novell's collaboration and productivity software (`www.novell.com/products/groupwise`).
- **ICQ:** (See the previous section of this chapter for information about creating an ICQ account.)
- **IRC:** As discussed in the "Instant Messaging Overview" section of this chapter, Internet Relay Chat (IRC) uses a multiuser chat room model rather than a standard, person-to-person instant messaging model. That's not a problem, just a different approach—after all, IRC preceded instant messaging. Because IRC is fundamentally different than IM, IRC is discussed separately. See the section of this chapter entitled "Using IRC" for more information.

- **MSN:** Microsoft stuff. (See the previous section of this chapter for information about creating a MSN account.)
- **MySpaceIM:** The instant messaging protocol used by MySpace Messenger (`www.myspace.com/myspaceim`). You send MySpaceIM messages using your MySpace e-mail address and password.
- **QQ:** The protocol and associated clients used by QQ (`www.qq.co.za/`), a fast, text and voice chat service that also supports sending SMS messages to cell phones.
- **SIMPLE:** The native protocol used by the Simple Instant Messenger, now known as SIM-IM (`http://sim-im.org`). SIM-IM is a free, open source, multi-protocol client that supports XMPP, Oscar, Yahoo Messenger, Microsoft Messenger (MSN), and other IM services that use the same protocols.
- **Sametime:** The protocol used by Lotus Instant Messaging (remember Lotus?), which was formerly known as "IBM Lotus Instant Messaging & Web Conferencing." Sametime is therefore used in many of IBM's commercial collaboration and messaging products (`www-306.ibm.com/software/lotus/sametime`).
- **XMPP:** (See the discussion of Jabber in the previous section of this chapter for information about creating a Jabber account.)
- **Yahoo:** (See the previous section of this chapter for information about creating a Yahoo account.)
- **Zephyr:** The protocol used by the Zephyr instant messaging system at MIT, long ago, as part of the same most excellent project (Project Athena) that brought you the X11 Window system. Anyone who thinks that AOL should really have a patent on instant messaging through its acquisitions of Mirabilis and ICQ should do a little bit of research and a whole lot of litigation before bending over and paying royalties. See `www-tech.mit.edu/V119/N19/history_of_athe.19f.html` and `http://web.mit.edu/zephyr/doc/usenix/zslides.PS` for today's history lesson.

After selecting the protocol/service where you have an instant messaging account, enter your Screen Name (the name of your IM account at a given service), Password, and (optionally) a local alias that you want to use for this account. If your computer is secure, you can click the Remember password checkbox, which will prevent you from having to enter your password each time that you start Pidgin, but which will also store your password in plain text in Pidgin's configuration file.

Next, you can optionally select the New Mail notifications checkbox to cause Pidgin to notify you whenever you have new e-mail (see Figure 13-2). You can also customize the icon that you want to associate with this account by selecting the Use this buddy icon… checkbox and clicking on the icon to display a dialog that lets you navigate to a directory where the custom icons that you want to use are located.

CAUTION **Pidgin's GAIM and Gaim predecessors showed distinct icons for users based on the protocol/service that they were using, but this default was dropped in Pidgin 2.0 because the developers felt that you were more interested in which people you were communicating with, rather than the protocol that they were using.**

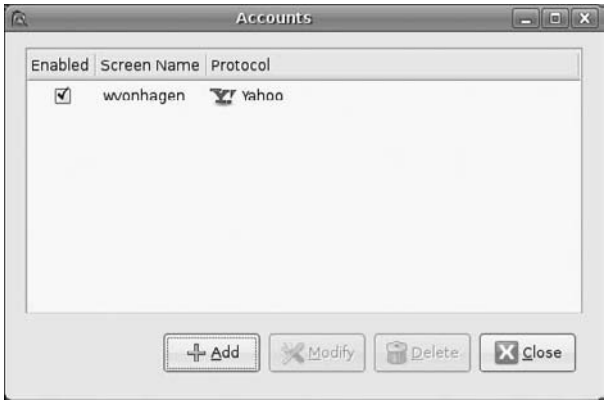
When configuring a Pidgin account, you can select the Advanced tab if you need to do things such as communicating with a specific server, use a specific port for communication, or want to use a proxy server so that your friends won't know your IP address. (The last can be useful for business-related instant messaging when you're working from the golf course, the local Starbucks, and so on.) The options that are provided on the Advanced tab are specific to the protocol/service that you're using. In most cases,

you won't need to set anything on this tab, but it's nice that Pidgin provides these sorts of advanced connectivity options if you do actually need them.

Once you've finished providing the details of your IM account, click Save to create the account in Pidgin. If this is the first time that you are configuring an IM account in Pidgin, Pidgin will display both the Buddy List window, shown in Figure 13-3, through which you can communicate with your IM friends, and the Accounts dialog, shown in Figure 13-2, from which you can add, modify, or delete IM accounts at any time. You can display the Accounts dialog at any time when you are running Pidgin by selecting the Accounts ➪ Manage menu option from the Buddy List window.

FIGURE 13-2

Pidgin's Accounts dialog



If you are done letting Pidgin know about your IM accounts, click Close in the Accounts window to close that window. Otherwise, you can click Add to define another IM account, or select an existing account and click Modify to change its settings, or click Delete to delete that account. Deleting an IM account only removes Pidgin's information about that account—the account still exists at the remote service.

TIP **Because Pidgin uses the Purple library for communication and configuration, all of your Pidgin account and buddy information is stored in the directory `~/.purple`. Specifically, all of your account information is stored in the file `~/.purple/accounts.xml`. If you want to completely purge this information from your computer and reconfigure Pidgin from scratch, you can simply delete this file or the entire directory if you want a truly clean slate.**

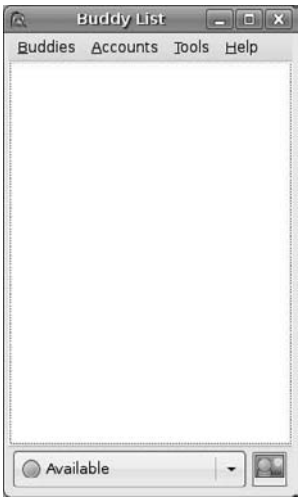
Once you've configured Pidgin, you're ready to start exchanging instant messages with your friends.

Using Pidgin

Pidgin's Buddy List dialog is the primary dialog through which you initiate instant messages and manage the list(s) of people with whom you will communicate frequently. Figure 13-3 shows an empty Pidgin Buddy List window, which is what you'll probably see the first time that you start Pidgin if you're new to instant messaging.

FIGURE 13-3

An empty Buddy List dialog in Pidgin



The next few sections discuss how to define IM contacts (in other words, how to add users to your Buddy List), how to organize your buddies, and how to extend Pidgin's basic capabilities by adding plug-ins. However, the first thing that you may want to do with Pidgin is simply to send someone an instant message—buddy lists and other conveniences make Pidgin more convenient and more fun, but sending instant messages is what IM is really all about!

Enter the screen name or an alias for the person that you want to send a message to, and click OK. A dialog like the one shown in Figure 13-5 displays on your screen.

To send an instant message to someone, regardless of whether they're shown in your Buddy List window, you can always select the Buddies ➪ New Instant Message menu option. (Once someone is in your Buddy List, you can send them an instant message by double-clicking on their name in the Buddy List.) If you've only defined a single instant messaging account, you'll see a dialog like the one shown in Figure 13-4. (If you've defined multiple instant messaging accounts, a similar dialog displays, in which you must identify both the person that you want to IM and the account that you want to use to do so.)

FIGURE 13-4

Initiating an instant message



FIGURE 13-5

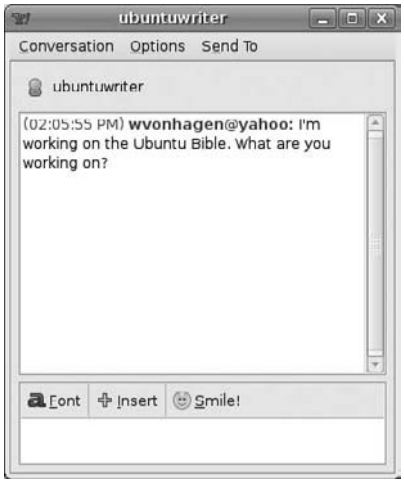
An instant message dialog



The heading of this window identifies the user that you are sending an instant message to, and the icon and name of that user is also listed above the conversation window. The conversation window is the top field on this dialog, and displays a record of your conversation as you send instant massages back and forth. The bottom portion of this dialog is where you type your messages. After typing a message to the user you identified, this dialog looks something like the one shown in Figure 13-6.

FIGURE 13-6

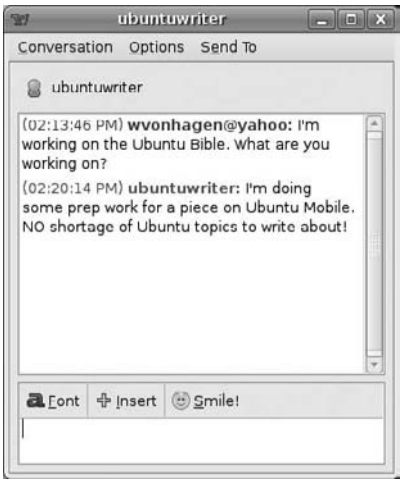
Sending a message in the conversation dialog



Each message in an instant message dialog is preceded by the screen name or alias for the user that is sending the message. In this case, although I'm sending the message using a Yahoo account named wvonhagen, you can see that my alias for this account is wvonhagen@yahoo, which is a convention that I use to make sure that I know which identity and service I'm using. If I hadn't defined an alias, my screen name would have been displayed. Once your friend responds, the instant message dialog will show your conversation to date, as shown in Figure 13-7.

FIGURE 13-7

Receiving a message in the instant message dialog



As you send messages back and forth, the last few messages will always be shown in the conversation window. When you're done communicating with your friend, click the Close box, select the Conversation ➪ Close menu option, or press Ctrl+W to close the conversation dialog.

Welcome to IM! The next few sections explain how to make using Pidgin a bit easier by defining buddies and how to do some other common tasks in Pidgin.

Connecting to Your IM Accounts in Pidgin

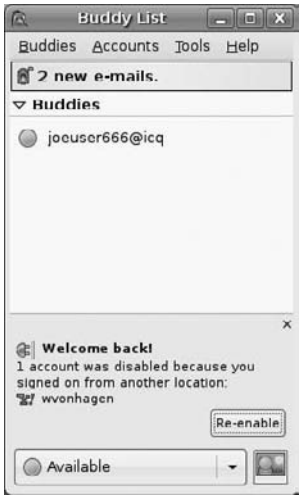
I typically have the same instant messaging account defined on multiple computer systems, such as my primary desktop system at home and my laptop, so that I can stay in touch with friends and family wherever I am. This is one of the great advantages of defining buddies (explained in the next section)—your buddy lists are stored by the service that you're using, so you see the same contact whenever you log in to a given service.

Most instant messaging protocols only allow you to be logged in from a single instant messaging client at a time. This makes sense, because otherwise every system where you are running an IM client would display chat and instant message windows, regardless of whether you were using that system or not. As mentioned in the previous section on defining IM accounts in Pidgin, you can set up those accounts to automatically log in whenever you start up Pidgin. This is usually a good idea for people with high-speed connections to the Internet, but the downside of this is that logging in from one computer system usually drops your

connection to the same IM account from any other computer system. When this occurs, Pidgin displays a message about this at the bottom of the Buddy List window, as shown in Figure 13-8.

FIGURE 13-8

A multiple login message in Pidgin



To re-enable your account immediately when you get back to that computer, you can click the Re-enable button in the message. You can click the X in the upper-right corner of this message if you simply want to close this dialog and re-enable the account later.

NOTE **Some IM services, like AIM/AOL, don't automatically drop other connections, but instead display a dialog that enables you to select the connection that you want to drop.**

If you need to re-enable an account manually, you can do so by selecting the Accounts menu's Enable Account submenu. This displays a list of all registered accounts that are not currently active—to re-enable one on your current computer system, simply select it from the drop-down list. If you are running a Pidgin client that supports that same account on any other machine, that account will be disconnected from the associated service on the other machine.

Adding Buddies in Pidgin

A buddy is the term for a person with whom you expect to communicate frequently via IM. As explained in the previous section, you don't actually have to add someone to your buddy list in order to send them an instant message, but you'll probably want to add your friends to your buddy list if you're going to communicate with them frequently (unless your memory is a lot better than mine is). When you first start using instant messaging, you may not have any predefined buddies, and your Buddy List may initially look like that shown in Figure 13-3.

To add a buddy, select the File ➪ Add Buddy menu option to display the dialog shown in Figure 13-9.

FIGURE 13-9

Defining a buddy in Pidgin



If you have defined multiple IM accounts in Pidgin, you'll first want to select the account that you want to use from the drop-down list at the top of this dialog. Next, enter the screen name of the person that you want to define as a buddy, an optional alias that you want to add them to, and (optionally, again) the group to which you want to add them. Groups provide a convenient way to organize your buddies so that you can easily recognize them as friends, family members, co-workers, and so on. This can be very useful as you accumulate more and more IM buddies—you may not immediately recognize the screen name of a person from whom you receive an instant message, so organizing them into groups will give you an immediate clue and should help clarify things. If you don't specify a group, your new buddy will be added to the default Buddies group.

After entering the screen name (and selecting an account, if necessary), click Add. The buddy that you have defined will immediately be listed in your Buddy List window.

Some instant messaging services, such as Yahoo, will automatically notify you when someone adds you to their buddy list. This displays a field at the bottom of the Buddy List window, as shown in Figure 13-10.

FIGURE 13-10

Authorizing someone to add you as a buddy in Pidgin



Clicking Deny displays an Authorization denied message dialog that enables you to send someone an IM that explains why you don't want them to add you to their buddy list. This is a message rather than a true rejection—you will still appear in their buddy list, and they can subsequently send you a message to clarify why they wanted to add you.

If you click Authorize, Pidgin conveniently displays a dialog like the one shown in Figure 13-9 with the screen name of the person who is trying to contact you filled in, enabling you to easily add them to your buddy list, too. After all, if you don't have a problem being on their buddy list, you might as well have them on your list, too. If you click Add in this dialog, they'll see a dialog like the one shown in Figure 13-10, which essentially just confirms that you OK'd their request and added them to their buddy list too.

As mentioned previously, once someone is in your buddy list, you can easily send him an instant message by double-clicking on his name in the Buddy List window.

Organizing Buddies in Pidgin

As you accumulate more and more IM contacts, your Buddy List window will fill up and you'll have to scroll further and further to actually find someone whose screen name you don't remember offhand. As mentioned in the previous section, you can use groups in the Buddy List window to organize your contacts. This usually feels like overkill when you only have a few buddies, but once you have a few zillion of them, you'll wish you'd organized them into groups from the beginning. Luckily, you can organize existing buddies into groups almost as easily as you can add them to groups when defining them in the first place.

By default, Pidgin provides the following groups into which you can organize your buddies:

- Buddies
- Co-workers
- Family
- Friends
- Recent Buddies
- Top Level

To change the group that a buddy is located in, right-click on that buddy, select the Move To menu item, and select the group to which you want to move that buddy.

You can create custom groups by selecting the Buddies ➪ Add Group command, which displays a dialog in which you can define the name of the new group. You can also rename an existing group by right-clicking on the name of that group and selecting the Rename command.

Viewing Buddies in Pidgin

By default, Pidgin only shows buddies who are online, which can be confusing after adding a buddy and then not seeing their name. You can change whether offline buddies are displayed in one of two ways:

- By selecting the Buddies ➪ Show command and selecting Offline Buddies. This will display all of your buddies, regardless of whether they are online or offline. Offline buddies will be shown as grayed-out entries with the word Offline shown below their identities. If you are using the default icons for this buddy, the standard green circle (online) icon will be replaced by a circle with an "x" in the middle when they are offline.
- By right-clicking on the name of a specific buddy in the Buddy List window and selecting Show when offline from the context-sensitive menu that displays. This buddy will be shown in your buddy list, regardless of their status or whether offline buddies are being displayed in general.

Using the Buddies ➪ Show ➪ Offline Buddies command can also be very handy if you aren't sure if you've added someone as a buddy or not, or when reorganizing your contacts into different groups. Selecting this command temporarily will enable you to see all of the buddies that you have defined so far, and will also enable you to move existing buddies to other groups, as explained in the previous section.

Logging Your Conversations

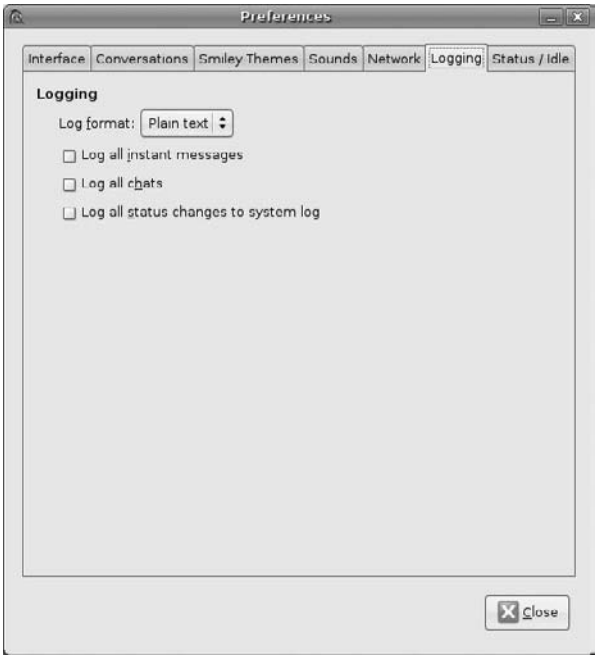
Depending on how compulsive you are, you may want to save the conversations that you have with certain people over IM. This can be important if you're using IM to get detailed information that you will want to refer to later, or simply if you want to remember what you said to someone the last time that you exchanged messages with them.

You can keep a log of a specific conversation by selecting the Options ➪ Log Messages menu option before initiating a conversation with someone else. Selecting this menu option logs all subsequent messages in a conversation, so this is only convenient if you know beforehand that you want to keep a record of your conversation. Selecting this command displays a message in your conversation window, confirming that logging has been enabled—the remote user doesn't see this message.

If you always want to keep a log of your conversations, you can configure Pidgin to do this by selecting the Tools ➪ Preferences menu option, and selecting the Logging tab in the Preferences dialog, shown in Figure 13-11.

FIGURE 13-11

Configuring logging preferences in Pidgin



Select the Log all instant messages checkbox to log all instant messages to a text file. If you want to log messages in HTML format instead of text, click the Log Format drop-down, and select HTML. You can also log all chats by selecting the associated checkbox. Once you are finished configuring logging in Pidgin, click Close to close the Preferences dialog.

If you don't want to automatically log all of your chats, you can always disable logging and simply use the Conversation ➪ Save As menu option in a chat window to save a specific chat to a file.

Exiting Pidgin

As mentioned earlier in this section, Pidgin creates an icon for itself in your panel's notification/system tray area when you start Pidgin. Closing the Buddy List window by selecting Buddies ➪ Quit only closes the Buddy List Window—it does not terminate Pidgin. After closing the Buddy List windows, Pidgin is still running and you can redisplay the Buddy List window at any time by left-clicking on Pidgin's icon in the notification/system tray portion of the panel. The fact that Pidgin is still running will cause a slight performance glitch if you try to restart Pidgin from the Applications menu when it is actually still running in the system tray. An entry that says "Starting Pidgin" will appear in the window list portion of the panel, but will eventually close without opening a Pidgin window on your screen.

TIP **If you do not want Pidgin to add an icon for itself in the system tray and continue running in the background once the Buddy List window is closed, select the Tools ➪ Preferences menu option to display the Preferences dialog, select the Interface tab, and select None from the Show system tray icon drop-down list. As a compromise, you can select On unread message to cause the notification area icon to continue to be displayed after closing the Buddy List window only if there are unread messages pending.**

To actually terminate Pidgin, you must right-click on the Pidgin icon in the notification portion of the GNOME panel (located at the far right of the top panel, by default) and select Quit from the content-sensitive menu that displays. If you select Quit while a Buddy List window is displayed on the screen, a Buddy List window will be displayed the next time that you start Pidgin. If you close the Buddy List window before selecting Quit from the Pidgin icon, whether the Buddy List window displays the next time that you start Pidgin depends on whether you've disabled the notification/system tray icon. If it is enabled (which is the default), only the notification/system tray icon is displayed the next time that you start Pidgin. If it is disabled, the Buddy List window will always display when you start Pidgin.

Getting More Information about Pidgin

As with all GNOME applications, you can get online help in multiple ways while using Pidgin:

- Pressing the F1 key or selecting the Help ➪ Online Help menu option takes you to the Pidgin FAQ on the Pidgin Web site, which may already contain the answer to your question
- Selecting the Help ➪ Get Help Online menu option takes you to the section of Ubuntu's Launchpad site that provides help for Pidgin in the release of Ubuntu that you're running. You can post a question there or review questions that others have already asked. This page also provides links to both the IRC channel where Ubuntu is discussed and to the Web-based Ubuntu forums. See the section of this chapter entitled "Using IRC" for more information about IRC.

As you might expect, the Pidgin Web site, `www.pidgin.im`, also provides a good deal of information about Pidgin, primarily in the Support and Development section that is located at `http://developer.pidgin.im`.

Instant Messaging on Kubuntu with Kopete

Kopete (`http://kopete.kde.org/`) is the multi-protocol instant messaging client that is provided with KDE versions 3 and 4. Earlier KDE instant messaging clients, such as Kit (`http://kitclient.sourceforge.net` or `www.hakubi.us/kit`) and Licq (`http://licq.sourceforge.net`), focused on AIM and ICQ compatibility, respectively, but Kopete focuses on supporting all of the popular instant messaging protocols that are supported today. Licq has subsequently added support for other protocols such as AIM and MSN, but Kopete is far and away your best choice for one-stop instant message shopping on KDE. Plus, it's installed as part of a default Kubuntu desktop installation, so there you go.

The name Kopete comes from the Chilean slang term "copete," which means something along the lines of "having a drink with friends." (The original author of Kopete, Duncan Mac-Vicar Prett, is from Chile.) This is nicely evocative of the spirit of instant messaging once you know about it, but somewhat less than obvious to those of us who aren't from Chile.

To start Kopete on a Kubuntu desktop system, select the Kopete icon from the Quick Launcher area in the panel, or select the K Menu ➪ Internet ➪ Kopete – Instant Messenger menu option. An icon for the Kopete application displays in the panel's System Tray area, and the main Kopete window displays. If this is the first time you've used Kopete, an account setup dialog displays that walks you through the process of configuring Kopete to know about any instant messaging account(s) that you already have, or setting one up.

Configuring Instant Messaging Accounts in Kopete

The first time that you start Kopete, a dialog appears that helps you configure Kopete to know about any instant messaging accounts that you have. Unlike most other instant messaging clients, Kopete's account configuration wizard will also help you create an instant messaging account if you don't already have one. Figure 13-12 shows the first account wizard dialog.

FIGURE 13-12

Selecting an IM protocol/service in Kopete



Figure 13-12 lists all of the instant messaging protocols/services that are supported by Kopete. If you already have an instant messaging account with one of these, select that protocol and click Next. If you do not have an instant messaging account yet, click the protocol/service that you want to join and click Next. The dialog shown in Figure 13-13 displays, in which you can configure Kopete to know about an existing account or register one so that you can then tell Kopete about it.

TIP **Of the types of accounts shown in Figure 13-12, Internet Relay Chat (IRC) is very different than standard IM. As discussed in the "Instant Messaging Overview" section of this chapter, Internet Relay Chat (IRC) uses a multiuser chat room model rather than a standard, person-to-person instant messaging model. That's not a problem, just a different approach—after all, IRC preceded instant messaging. Because IRC is fundamentally different than IM, IRC is discussed separately. See the section of this chapter entitled "Using IRC" for more information.**

FIGURE 13-13

Defining or creating an IM account in Kopete



The Account Information dialog shown in Figure 13-13 is customized based on the protocol/service that you selected in the preceding dialog. Figure 13-13 shows the Account Information dialog that you see after selecting the AIM protocol/service in Figure 13-12. Other screens are very similar.

If you do not yet have an account with the protocol/service that you selected in the dialog shown in Figure 13-12, click Register New Account in the lower half of this dialog. This will open the registration page for that protocol/service in Konqueror, so that you can create an account. This is an amazingly friendly feature—no wonder that KDE and the KDE developers always get kudos for user-friendliness!

Enter the screen name for your new or existing IM account in the Screen Name field. If you don't want to be prompted for your password for this account each time you start Kopete, select the Remember password checkbox, which activates the Password field, in which you can enter your password. Kopete will

store your password in the KWallet credentials management application, so you don't have to worry about your password being stored in plain text somewhere on your system. You can optionally select Exclude from connect all to prevent this account from automatically logging in when you start Kopete, which you may want to do if your have a low bandwidth connection to the Internet or if you simply want to manually choose which IM accounts you connect with each time that you start Kopete. (For example, you may want to prevent work-related IM accounts from automatically connecting whenever you start Kopete.) You can also optionally select the Exclude from Global Identity checkbox if you don't want this IM account to be consolidated into the global identity that Kopete provides to simplify updating your personal information. (More about that later.)

After entering your account information and making any optional selections, click Next to proceed. The dialog shown in Figure 13-14 displays.

FIGURE 13-14

Kopete's final account definition dialog



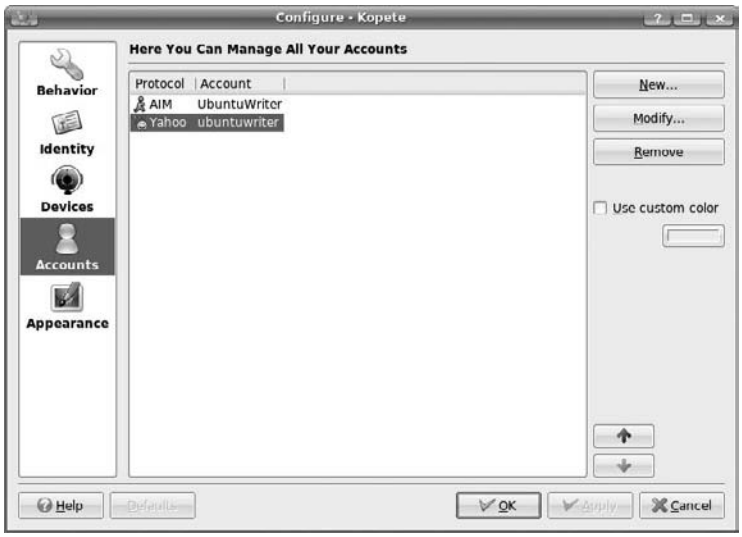
This final dialog optionally enables you to use a custom color to identify the account you've just defined. By default, the Connect now checkbox is selected, which will cause you to connect to the protocol/service associated with the IM account you've just defined. Click Finish to close this dialog and connect to that service if you leave that checkbox selected.

You can subsequently add (or create) new IM accounts for use in Kopete by selecting the Kopete window's Settings ➪ Configure menu option and selecting the Accounts icon to display the pane shown in Figure 13-15.

This pane enables you to create new accounts, modify an existing account (by selecting that account and clicking Modify), or remove an existing account (by selecting that account and clicking Remove.

FIGURE 13-15

Using Kopete's account configuration pane



Using Kopete

A dialog like the one shown in Figure 13-16 displays when you first start Kopete. This window, generally referred to as the Contact List window, shows any buddies, custom groups, and default groups associated with the protocol(s)/service(s) that you're using for instant messaging.

FIGURE 13-16

The default Kopete dialog



To send an instant message to any existing buddy, left-click on the name of that buddy in the contact list. A chat window displays, as shown in Figure 13-17.

FIGURE 13-17

Sending an Instant Message in Kopete



To display this window, I clicked on my ubuntuwriter buddy in the Kopete window. As you can see from Figure 13-16, we've already exchanged a round of messages, and I have a response waiting (in the field at the bottom of the window) to be sent as soon as I press Return.

You can communicate with as many buddies as you want, and simultaneously. Selecting another buddy from the Kopete window causes another tab to open in the Chat Window, as shown in Figure 13-18. The title of the window reflects the currently selected tab, and in this case indicates that my friend is away from his computer. I can still send him an IM, which he'll see as soon as he gets back.

As you can see, sending and receiving instant messages in Kopete is easy. Kopete provides an attractive interface for instant messaging. As we'll see in the next few sections, it also makes it easy to do standard instant messaging tasks, and has a few special capabilities that make it both unique and attractive.

Connecting to Your IM Accounts in Kopete

I typically have the same instant messaging account defined on multiple computer systems, such as my primary desktop system at home and my laptop, so that I can stay in touch with friends and family wherever I am. This is one of the great advantages of defining buddies (explained in the next section)—your buddy lists are stored by the service that you're using, so you see the same contact whenever you log in to a given service.

FIGURE 13-18

Having multiple conversations at one time



Most instant messaging protocols only allow you to be logged in from a single instant messaging client at a time. Logging in from one computer system usually drops your connection to an IM service from any other computer system or allows you to select which client to disconnect. This makes sense because otherwise every system where you are running an IM client would display chat and instant message windows, regardless of whether you were using that system or not. To cause Kopete to automatically log you in on all of the accounts that you have defined (which will log you out on any other clients where the same account is enabled), select the Settings ➪ Configure menu option, select the Behavior icon in the left-hand toolbar, and check the Connect automatically at startup checkbox, as shown in Figure 13-19.

You can exclude specific accounts from automatically connecting regardless of whether the Connect automatically option is enabled by selecting the Exclude from connect all option when you create that account, as explained earlier in the section "Configuring Instant Messaging Accounts in Kopete." However, if you do so, you'll need to manually connect those accounts to their associated protocols/services before you can exchange instant messages with buddies from that service.

You can normally only exchange messages with buddies who are online, and then only if you're connected to that service. Trying to send a message to a buddy when you're not connected to a service displays a message to that effect. The icons in the bottom-right corner of the Kopete window represent each of the accounts that you have defined in Kopete. To manually connect to a service using one of these accounts, right-click on that icon and select the Online menu option, as shown in Figure 13-20.

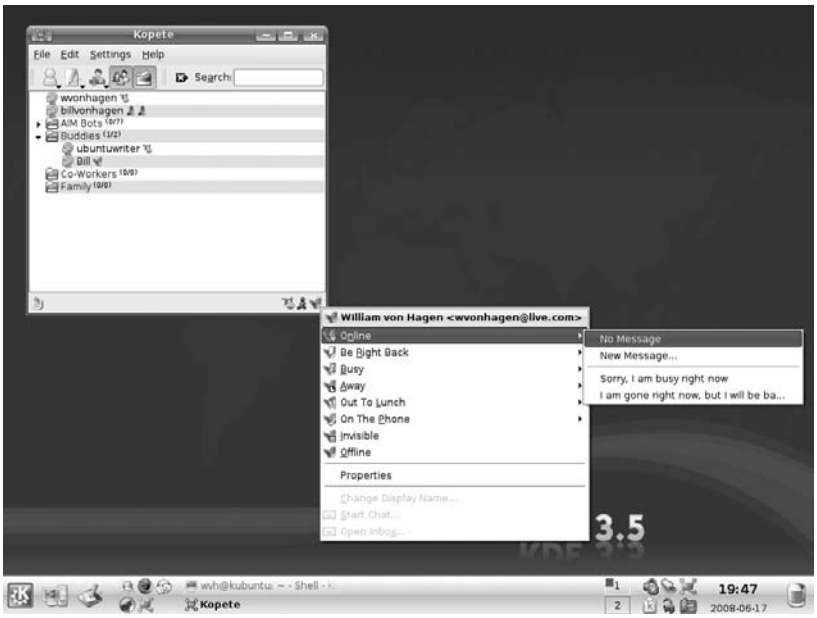
FIGURE 13-19

Configuring Kopete to automatically connect clients



FIGURE 13-20

Manually connecting to an IM service



Some services enable you to send a message when you connect to that service, which will be sent to all of your buddies. I generally prefer to simply connect, so I typically select the No Message option when it's available (as shown in Figure 13-20).

You can also use the icons in the lower-right corner of the Kopete window to manually disconnect from a service if marking yourself as "Busy" (which shows you as being busy, optionally displaying a custom message) or "Invisible" (which means that no one can detect that you're logged in) simply isn't good enough.

One interesting feature of Kopete is its support for a global identity. After all, you may have multiple instant messaging accounts with multiple services, but they're all really you, right? Kopete enables you to configure and use a global identity that you can link to Kontact's KAddressbook module, enabling you to update the identity information for all associated instant messaging accounts at one time. This feature is off by default, but you can activate a global identity by selecting the Settings ➪ Configure menu item, and selecting Identity from the icon toolbar at left to display the pane shown in Figure 13-21.

FIGURE 13-21

Configuring a global identity in Kopete



Once this dialog displays, you can then check Enable global identity and configure a default identity or create a new one to use. Each identity enables you to determine which of your accounts will supply the nickname for all other IM accounts when that identity is being used. You can also use the Photo tab to associate a photo with each identity, link that identity to a KAddressBook contact using the Address Book Link tab, and so on.

Adding Buddies in Kopete

A buddy is the term for a person with whom you expect to communicate frequently via IM. As explained in the previous section, you don't actually have to add someone to your buddy list in order to send them an instant message, but you'll probably want to add your friends to your buddy list if you're going to communicate with them frequently (unless your memory is a lot better than mine). When you first start using instant messaging, you may not have any predefined buddies, and your Buddy List may initially be pretty empty.

There are two ways of defining buddies—either by defining them manually, or by adding people to your buddy list after they contact you. To add a buddy manually, right-click in the Kopete window and select the Add Contact menu item, as shown in Figure 13-22.

FIGURE 13-22

Adding a contact manually in Kopete



Selecting this menu item causes a pop-up menu to display, from which you can select the account for which you want to define a buddy. This displays a dialog like the one shown in Figure 13-23.

FIGURE 13-23

Specifying buddy/contact details in Kopete



This dialog enables you to define the name of the buddy/contact that you want to add, the group to which you want to add that buddy, and so on. Groups provide a convenient way to organize your buddies so that you can easily recognize them as friends, family members, co-workers, and so on. This can be very useful as you accumulate more and more IM buddies—you may not immediately recognize the screen name of a person from whom you receive an instant message, so organizing them into groups will give you an immediate clue and should help clarify things. If you don't specify a group, your new buddy will be added as a top-level contact. Once you're done, click OK to save the new contact.

Adding buddies yourself is fine, as long as you can think of all of your friends. In many cases, one of your friends will find your IM account and contact you. In this case, Kopete displays a message balloon about the new message from the Kopete icon in the system tray, and also adds an entry for that contact to Not in your contacts list in the main Kopete window, as shown in Figure 13-24.

FIGURE 13-24

A message from a new contact in Kopete



Assuming that you recognize the new contact, perhaps after exchanging instant messages with them, you can add them to your buddy/contact list by right-clicking on their name in the Kopete window, which displays the pop-up menu shown in Figure 13-25.

Select the Add to Your Contact List menu option to permanently add them to the list of buddies/contact for an IM account. If the Not in your contacts list section of the main Kopete window is now empty, that section is removed.

Organizing Buddies in Kopete

As you accumulate more and more IM contacts, your main Kopete window will fill up and you'll have to scroll further and further to actually find someone whose screen name you don't remember offhand. As mentioned in the previous section, you can use groups in the Kopete window to organize your contacts. This usually feels like overkill when you have only a few buddies, but once you have a few zillion of them, you'll wish you'd organized them into groups from the beginning. Fortunately, you can organize existing buddies into groups almost as easily as you can add them to groups when defining them in the first place.

Many instant messaging services and clients provide some number of top-level groups as a starting point for your organizational efforts. By default, Kopete does not provide any groups of its own, leaving it up

to you to create the groups that are meaningful to you rather than asking you to shoehorn your buddies/contacts into a set of preconceived groups.

Kopete makes it very easy for you to create your own groups. To do this, select the File ➪ Create New Group menu option. A dialog like the one shown in Figure 13-26 displays.

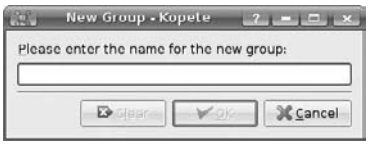
FIGURE 13-25

Adding a new contact in response to an IM



FIGURE 13-26

Creating a group in Kopete



Enter the name of the group that you want to create, and click OK.

To move an existing contact to a new group, right-click on the name of that contact, select the Groups menu item, select the Move To command from the menu that displays, and select the name of the group to which you want to move that contact from the list of groups that appears.

Kopete also supports having the same contact in multiple groups. This can be incredibly useful when you interact with a single person in a number of different ways—perhaps as a co-worker or friend, and through some shared special interest. You can create groups for each of these sets of related users, copy appropriate users into each of them, and use that group as a way of locating different contacts in different contexts. Copying an existing contact to another group is very similar to moving a contact—to copy an existing contact to another group, right-click on the name of that contact, select the Groups menu item, and select the Copy To command from the menu that displays. Select the name of the group to which you want to copy that contact from the list of groups that appears. If the contact is a top-level contact, it will actually be moved to that group; otherwise, an entry for that contact will be copied to the selected group.

Groups not only provide a way of organizing contacts, but also provide a number of conveniences when communicating with the contacts in a group. Right-clicking on a group displays the pop-up menu shown in Figure 13-27, which provides access to some of the more common group-related tasks in Kopete.

FIGURE 13-27

Group-related actions in Kopete



Of the tasks shown in the pop-up menu in Figure 13-27, I find the Send Messages to Group command to be very useful. Because Kopete enables you to list the same contacts in multiple groups, being able to send a single message to all of the members of a group can be very convenient because you can put certain contacts in multiple groups. For example, people that I've played in bands with are typically simply buddies of mine, but I have copied some of them into a group called Punk Music for fans of that particular genre. I can then use the Send Messages to group function to broadcast messages of special interest to that group, such as when I want to report the demise of yet another aging music legend.

Kopete also enables you to set custom notifications for each group. Normally, you're notified if incoming instant messages arrive, but you can use a group's Properties dialog to define custom notifications that should be performed when events such as incoming messages or contact status changes occur. To set the properties for a group, right-click on the name of the group and select Properties. Once the Properties dialog displays, click the Custom Notifications tab to display the pane shown in Figure 13-28.

FIGURE 13-28

Setting custom notifications for a group



Figure 13-28 shows the menu of events with which you can associate special actions, such as playing a sound, displaying a message, starting a chat, and so on. You can even select the Suppress standard notifications checkbox to suppress standard events, such as pop-up messages, that typically occur when events such as new messages occur. This is discussed in the next section.

Customizing Incoming Message Events

When instant messaging first appeared on Linux systems, it was often irritating. Incoming messages in most clients automatically raised a messaging window when an incoming message arrived, and the new window stole the focus from whatever other window you were working in. I found this so irritating that I rarely used instant messaging. Nowadays, things don't work that way by default, although you can still activate this behavior if you want.

By default, new instant messages in Kopete cause a balloon notification to display on your screen, and cause the Kopete icon in the system tray to flash/rotate. This seems like overkill to me, so I like to customize what happens when new messages arrive—which you can easily do from Kopete's Configure dialog. To customize how Kopete reacts to system events, select the Settings ➪ Configure menu option, select the Behavior icon in the left-hand icon bar, and select the Events tab to display the pane shown in Figure 13-29.

FIGURE 13-29

The Events tab of Kopete's Configure dialog



This pane provides total control over what happens when an incoming message is received. Deselecting the Show bubble checkbox prevents a speech balloon (or bubble, in Kopete terms) from being displayed whenever an incoming message is received. By default, the bubble displays on your screen until you either click View (to open a chat window showing that message) or Ignore (to close the bubble). If you still like the bubble but want to fine-tune how the speech balloon (or bubble, in Kopete terms) displays, you can

select the Close bubble automatically checkbox and specify the period of time that the bubble will display on your screen. Similarly, if you don't like the flashing Kopete icon in the system tray when an incoming message has been received, you can deselect the Flash system tray option. And finally, if you want the Kopete window to pop to the top and steal the cursor from whatever application you're currently using, you can select the Raise window on incoming message option in the Miscellaneous section.

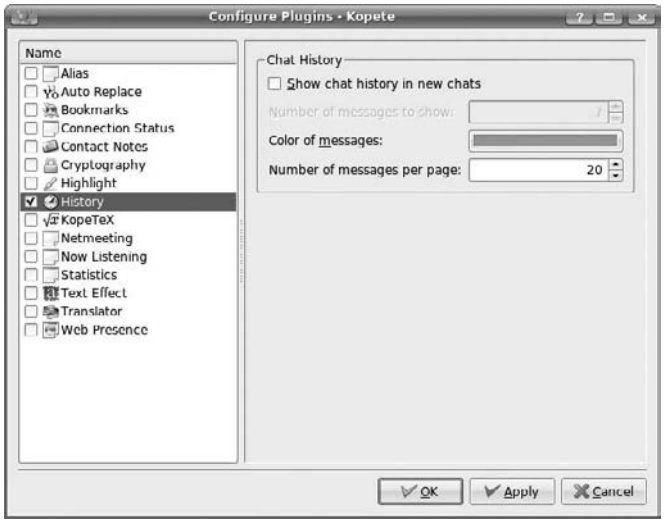
Logging Your Conversations

Depending on how compulsive you are, you may want to save the conversations that you have with certain people over IM. This can be important if you're using IM to get detailed information that you will want to refer to later, or simply if you want to remember what you said to someone the last time that you exchanged messages.

Kopete's logging is controlled by its History plug-in. Plug-ins provide a flexible mechanism for extending applications such as Kopete by encapsulating functionality in small packages that are separate from the main Kopete executable. Adding plug-ins enables you to add new capabilities to applications such as Kopete without having to recompile the application or install a newer version from a package. Kopete's History plug-in is normally enabled by default, but you can verify whether it is active (and activate it, if necessary) by selecting the Settings ➪ Configure Plugins menu option to display the dialog shown in Figure 13-30.

FIGURE 13-30

Configuring Kopete plug-ins



If no checkmark is displayed to the left of the entry for the History plug-in, select that checkbox and click OK to enable the plug-in. You can then see the history of your chats with any user by right-clicking the contact for that user and selecting the View History menu option. This displays a window like the one shown in Figure 13-31, which enables you to select between multiple chats that you've had with that user.

FIGURE 13-31

Viewing your chat history with a given contact



You can cut and paste the information from this window into any other application. The logs used by the history mechanism are stored in XML format in subdirectories of the directory `~/.kde/share/apps/kopete/logs`. These subdirectories are organized by protocol (AIMProtocol, YahooProtocol, and so on), with chat logs for your exchanges with different users in time-stamped files in these subdirectories.

If you don't want to automatically log all of your chats, you can always disable the history mechanism and manually save chats by selecting the Chat ➪ Save menu option to save a specific chat as a text file.

Exiting Kopete

As mentioned earlier in this section, Kopete creates an icon for itself in the kicker's system tray area when you start Kopete. Closing the Kopete window by using the Close box in the upper-right corner does not actually terminate Kopete—you can redisplay the Kopete window at any time by left-clicking on Kopete's icon in the system tray. To actually close Kopete, select the File ➪ Quit menu option, or right-click on the Kopete icon in the system tray and select the Quit command from the context-sensitive menu.

TIP **If you do not want Kopete to add an icon for itself in the notification area and continue running in the background once the Kopete window is closed, select the Settings ➪ Configure menu option to display the Configure dialog. If it is not already selected, select the Behavior icon in the left-hand icon bar and click the general tab. De-select the Show system tray icon option and click OK to prevent Kopete from creating a system tray icon the next time that you start Kopete. This will also close any system tray icon for Kopete that is currently shown.**

Getting More Information about Kopete

Kopete is a powerful and flexible application; exploring its every nuance could potentially require its own book. The following are some additional sources of information about Kopete:

- **Kopete Handbook:** Pressing the F1 function key while running Kopete displays the Kopete Handbook, which is part of KDE's online help. Although it hasn't been updated recently (the last update was at the end of 2006), the Kopete Handbook is still quite relevant because Kopete hasn't required significant changes recently. After all, why fix something if it isn't broken?
- **Kopete Web site:** (`http://kopete.kde.org/`) As you'd expect from the Kopete mothership, this site includes tons of useful information, including the Kopete FAQ, an online version of the Kopete handbook, a link to the Kopete Wiki, a link to the Kopete developer's mailing list, screenshots for multiple versions of Kopete, the Kopete roadmap (to see where Kopete is headed), the current Kopete bug list, and much more.
- Kopete site on SourceForge: (`http://sourceforge.net/projects/kopete/`) Provides access to the Kopete source code and hosts the Kopete forum (`http://sourceforge.net/forum/forum.php?forum_id=124733`), which is a valuable source of information and a good place to ask questions.

Hundreds of other sites and online articles provide additional information about Kopete—a quick Web search through search sites such as Clusty.com or Google.com can often help answer those tricky "how do I" questions.

Using IRC

As discussed in the "Instant Messaging Overview" section of this chapter, Internet Relay Chat (IRC) was the first cross-platform mechanism for real-time Internet communications. As mentioned earlier, IRC is a client/server. IRC users connect to an IRC server using a client application that runs on their systems. IRC servers are typically themselves organized into networks of servers that share information with each other. Hundreds, if not thousands, of IRC servers are publicly available and typically do not require a password in order to access them. (See `www.irchelp.org/irchelp/networks` for a good starting list of available IRC servers and networks.) The discussion areas on IRC are referred to as channels. Channel names are prefixed by a hash mark (#)—one of the most popular IRC channels for readers of this book should be the #ubuntu channel, which is a support channel for Ubuntu, a great place to ask Ubuntu-related questions at any level, and a generally useful resource. Most channels are accessible by anyone, but they can also be password-protected to provide private areas for communication between sets of users who know the password.

When you first connect to an IRC server, you define a nickname (generally referred to simply as a "nick"), which is how you are identified on any channel that you join. Once connected, you can use commands

prefaced with a slash (/) character to interact with the server itself. The following are common IRC commands:

- `/dcc chat nickname:` Request a direct, user-to-user chat session with the user who has the specified nickname.
- `/help`: Summarizes available IRC server commands. You can follow the `/help` command with the name of another command (without the `/`) to get specific information about that command.
- `/join #channel`: Enables you to connect to a specified channel. If the specified channel does not already exist, it will usually be created with you as its operator, but some IRC servers limit channel creation to specific users in order to cut down on the noise (and to prevent typos from automatically creating channels).
- `/leave #channel`: Leave the specified channel.
- `/list`: Displays a list of available channels.
- `/quit`: Drop your connection to IRC, displaying a farewell ("signoff") message.
- `/who #channel`: Lists the users that are currently connected to a specific channel.

Although the first IRC clients were for UNIX and UNIX-like systems, IRC clients have been available for other operating systems since shortly after the birth of IRC. Mac OS, Mac OS X, and Microsoft Windows systems all offer a wide selection of IRC clients. On the Windows platform, mIRC (`www.mirc.com/`) is probably the most popular IRC client, although many others are available. The Miranda IM client for Windows (`www.miranda-im.org`) also supports IRC, as does Trillian (`www.ceruleanstudios.com`). Macintosh users typically use IRC clients such as ircle (`www.ircle.com/`), which has been popular on Macintosh systems forever, Adium X (`www.adiumx.com`), which provides both IM and IRC access, and Colloquy (`http://colloquy.info`), which is impressive on both the OS X desktop and on your iPhone.

The next two sections discuss how to connect to IRC using the Pidgin and Kopete instant messaging clients—they aren't called multi-protocol clients for nothing! A third section provides tips for getting more information about IRC.

Using IRC in Pidgin

As discussed in the section "Instant Messaging on Ubuntu with Pidgin," IRC is one of the messaging protocols that is built into Pidgin. Because most IRC servers are open, you can create an IRC account and connect to an IRC server without having to pre-register, but you must create a Pidgin account to do this. You create your Pidgin IRC account like any other account, by selecting the Accounts ➪ Manage menu option and selecting the IRC protocol, as shown in Figure 13-32.

Your screen name is your IRC nickname, which can be whatever you want it to be, but which must be unique on the IRC server that you're connecting to. By default, Pidgin IRC accounts connect you to the IRC server `irc.ubuntu.com` (part of the freenode IRC network), but you can change the name of the server if you want to connect to another IRC server for some reason.

Once you save your Pidgin IRC account, you'll see the standard Buddy List dialog.

Figure 13-33 shows the chat window that you'll see after connecting to the `irc.ubuntu.com` IRC server.

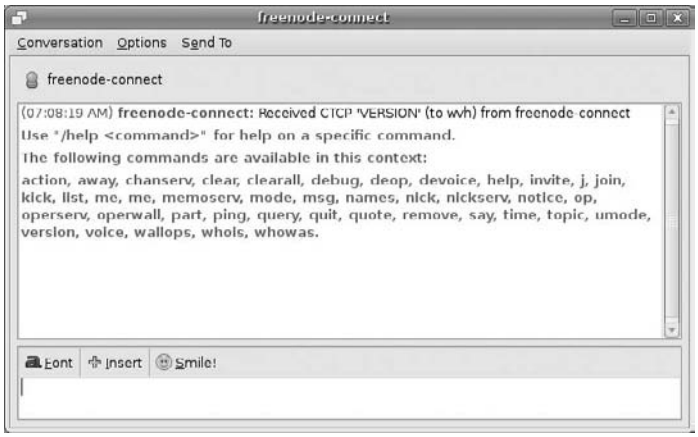
FIGURE 13-32

Creating an IRC account in Pidgin



FIGURE 13-33

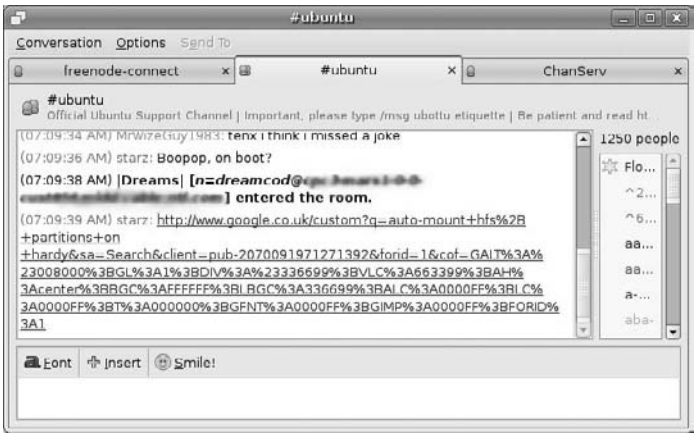
An initial connection to an IRC server



Once connected to an IRC server, you'll want to either use the `/list` command to see available channels, or use the `/join` command to join a specific chat. Joining a channel opens a new tab for communication on that channel, and also opens a ChanServ (channel server) tab in the background, which displays channel registration, administration, and general information messages for channels that have registered with the administrative ChanServ service. Figure 13-34 shows a connection to the #ubuntu channel after typing `/join #ubuntu` in the text window at the bottom of the Pidgin chat window.

FIGURE 13-34

Joining the #ubuntu channel



At this point, you can enter other IRC commands in the text window at the bottom of the pidgin screen, or simply enter text to send it to the channel. When you're done asking questions, answering questions, or simply watching the flow of communications, you can leave the channel by using the `/leave` command to leave a channel, or the `/quit` command to close your connection to the IRC server.

NOTE **If you find IRC interesting, useful, or perhaps even both, you may want to investigate other Linux IRC clients. Although Pidgin is fine for my purposes, the XChat program is a UNIX/Linux IRC client that has been around for a long time, is quite popular, and makes it very easy to connect to multiple channels and servers at the same time. XChat uses GTK, so it is well-suited for use on GNOME-based Ubuntu desktop systems—or on Kubuntu systems if you're willing to install its GTK dependencies. XChat isn't installed by default, but is available in the default Ubuntu repositories. The latest version of XChat for Ubuntu systems is located in the xchat-gnome package, which you'll have to install using Synaptic (or the adept_installer on Kubuntu systems), aptitude, or apt-get.**

Using IRC in Kopete

As discussed in the section "Instant Messaging on Kubuntu with Kopete," IRC is one of the messaging protocols that is built into Kopete. Because most IRC servers are open, you can create an IRC account and connect to an IRC server without having to pre-register, but you must create a Kopete account to do this. You create your Kopete IRC account like any other account, by selecting the Tools ➪ Configure menu option, selecting the Accounts icon in the left-hand toolbar, and clicking New. The first screen of the Add Account Wizard displays, from which you can select the IRC protocol. The screen shown in Figure 13-35 displays.

Entering the nickname that you want to use is the only requirement for connecting to IRC through Kopete, although you may want to specify an alternate nickname if your first choice is already in use on the server that you connect to. You can also enter a Real Name if you want IRC users to be able to see this name if they request information about you using the `/whois` command.

Next, if you want to connect to a specific IRC network, you may want to select the Connection tab and use the Network drop-down list to select the network that you want to join, as shown in Figure 13-36.

FIGURE 13-35

Creating an IRC account in Kopete

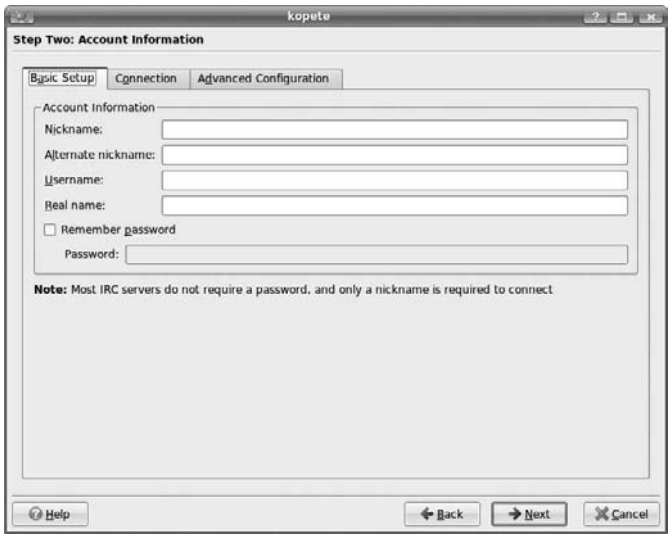
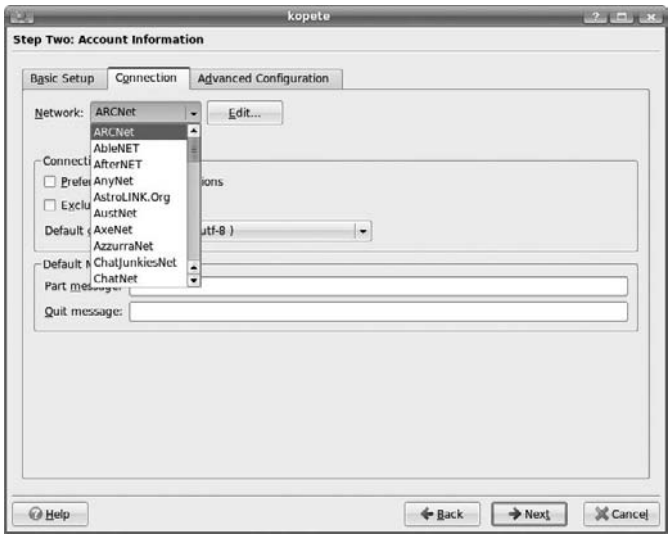


FIGURE 13-36

Specifying an IRC network to connect to



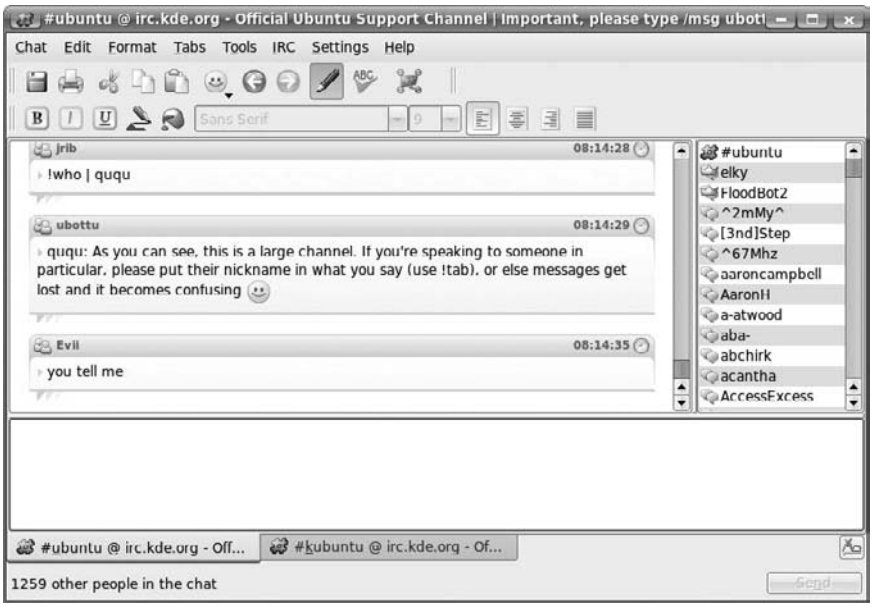
My personal favorite IRC network is freenode, which appears near the bottom of the list because its name is specified in lowercase, but you can select any network that you want to use. If you do not specify a network to connect to, the default IRC ARCNet network is used. Once you've selected a network, you can also click Edit if you want to identify a specific server on that network to connect to.

After completing your network settings, click Next to proceed to the final account wizard dialog, where you can deselect the Connect Now checkbox if you don't want to connect immediately and associate a specific color with this Kopete account. Click Finish to exit the account wizard. Click OK to close Kopete's configuration dialog.

Next, assuming that you left the Connect Now checkbox selected when you created your IRC account, you can connect to a specific channel by left-clicking the icon for your IRC account in the lower-right corner of the Kopete window and selecting the Join Channel command. A dialog displays, in which you identify the name of the channel that you want to join. For example, to join the Kubuntu support channel, you would enter #kubuntu and click OK. Similarly, you could enter #ubuntu to join the Ubuntu support channel. A Kopete chat window like the one shown in Figure 13-37 displays.

FIGURE 13-37

The #ubuntu channel in Kopete



Once connected to an IRC channel, you can send messages to the channel by typing them in the text area at the bottom of the Kopete window. Kopete provides the Chat ➪ Commands drop-down menu to enable you to send IRC commands, with the Close command leaving the current channel and closing its tab. Each channel that you connect to using the Chat ➪ Commands ➪ Join command opens in a separate tab in your IRC chat window.

Getting More Information about IRC

Regardless of how well it fits the IM model, IRC is still popular, quite fun, and occasionally even useful. When you first connect to IRC, you'll probably be doing so in order to connect to a specific channel, like #ubuntu. However, channels such as #new2irc, #newbies, or #newuser are also interesting starting points. If you'd like general information about IRC or have some specific questions, you may want to explore #irchelp. For information about available channels, use `/list` on the IRC server that you're connected to, or use a Web-based resource such as `www.irchelp.org/irchelp/chanlist`.

The following are some great sources of additional information about IRC, its culture, IRC servers, other IRC sites, and so on:

- **IRC.org:** (`www.irc.org`) Known as the original home of IRC, this site provides news, help, IRC history and humor, information about ongoing developments in IRC, and a huge collection of useful links to many other IRC-related sites.
- **IRChelp.org:** (`www.irchelp.org`) The Internet Relay Chat Help site provides a tremendous collection of IRC FAQs, primers, HOWTO guides, helper scripts, lists of active servers and supported channels, and so on.
- **Wikipedia:** (`http://en.wikipedia.org/wiki/Internet_Relay_Chat`) In addition to the standard "online encyclopedia" info, the Wikipedia article on IRC provides good links to many other IRC-related sites—or at least it did at the time of this writing.

Other Instant Messaging Packages

Pidgin and Kopete provide all of the functionality that almost anyone needs in a multi-protocol IM client. However, Ubuntu is a Linux distribution, so you might suspect that there are a huge number of other IM applications out there. Some of my favorites are the following (listed in alphabetical order so as not to play favorites):

- **Google Talk:** (`www.google.com/talk`) Google Talk is a Windows and Web-based client for Google's chat service, or any other XMPP-based chat service.
- **Licq:** (`www.licq.org`) Licq is an ICQ client that uses the Qt desktop widgets. Still actively under development. Plug-ins are available to enable Licq users to communicate with AIM and MSN users.
- **Meebo:** (`www.meebo.com`) Meebo is a Web-based instant messaging service that supports its own accounts as well as providing access to other IM services that Meebo supports, such as AIM, GoogleTalk, ICQ, jabber, MSN, and Yahoo. A Firefox plug-in for Meebo is also available.
- **Psi:** (`http://psi-im.org`, also in Ubuntu repositories) Psi is a free Jabber/XMPP Client that is available for Linux, Macintosh, and Windows systems and uses Qt for its graphical interface. Psi is a great Jabber client, but you can also use Psi to connect to other IM services by taking advantage of gateway/transport services such as those offered by the `im.flosoft.biz` Jabber server. Figure 13-38 shows Psi with a few gateway connections defined.
- **SamePlace:** (`http://sameplace.cc`) SamePlace is a Web-based XMPP instant messaging system that provides Firefox, Flock, and Thunderbird plug-ins. You can also take advantage of its gateways to other services (known as transports, or transport services) to use SamePlace to communicate with other supported IM services.
- **ScatterChat:** (`http://scatterchat.en.softonic.com`) is an encrypted, anonymous instant messaging application, and is designed to provide a lifeboat of privacy in today's sea

of online communication. ScatterChat is a product of the Hacktivismo project (`http://en.wikipedia.org/wiki/Hacktivismo`), which views access to information as a basic right. And right they are.

FIGURE 13-38

The Psi Jabber client and transports to other services



Many other IM clients are available in the open source universe, and I can't list them all here. Heck, I probably don't even know them all because I'm sure that one was just written yesterday. However, between the excellent and rich capabilities of Pidgin and Kopete and the packages listed in this section, I'm sure that you can find the instant messaging client that's just right for you.

Resolving IM Communication Problems

Both Pidgin and Kopete are popular applications that are used daily by many people. If you're having problems sending or receiving instant messages with either of these applications, the most common source of communication problems is that a firewall is blocking your connection to the instant messaging service that you're trying to reach. If you are experiencing communication problems and are running a firewall, make sure that the appropriate port(s) from the following list are not being blocked, redirected, or being used for something else:

- AIM/ICQ—port 5190
- Google Talk, Jabber, and XMPP in general—port 5222
- IRC—port 6667
- MSN—ports 443 (SSL) and 1863
- Yahoo—port 5050

Blocking Internet-wide access to these ports is quite common at many businesses because of the possibility that randoms on the Net could eavesdrop on business-related IM communications. Many companies offer enterprise IM solutions, but that's certainly outside the scope of this book.

If you are trying to use instant messaging on your personal computer and are not running a firewall, you may want to check with your ISP to make sure that they are not blocking these ports. That would be incredibly antisocial, but not completely surprising.

Summary

Instant messaging is a popular interpersonal communication mechanism for most Internet-connected people nowadays. This chapter provided some background and historical information on instant messaging, and explained how to obtain an instant messaging account with various popular instant messaging services. The next two sections provided detailed configuration and usage information for the instant messaging clients that are installed on Ubuntu and Kubuntu desktop systems by default. The chapter also discussed Internet Relay Chat—a related multi-person communication mechanism—highlighted other instant messaging clients that are available for Linux systems, and concluded by providing some trouble-shooting information that may help you if you are having problems communicating with your instant messaging service of choice.

Chapter 14

Using Command-Line Tools

Because Linux has its conceptual roots in the older UNIX operating system, many Linux applications are executed from a command line. A command line is the traditional interface found on older computer systems that may not have used the high-resolution, graphically oriented monitors that most people expect today. In the command-line model, the system runs a program, known as a *command-line interpreter*, which does just what its name suggests: It reads the commands that you type, locates the appropriate application on your system, and executes that application for you as instructed based on what you've typed. Once the command completes, the command interpreter displays a sequence of characters, known as a *prompt*, signifying that it is ready to accept another command.

Linux systems aren't reliant on the command-line approach simply out of historical interest—when a Linux system boots, it needs to be able to run several commands long before the graphical interface is available. On Linux systems, the graphical interface itself is started by a command-line utility, after which (of course) fancy graphics are available, expected, and used.

For most people, simply displaying a command-line prompt on a fancy graphical monitor powered by a multi-megabyte graphics card would be a waste of good hardware. However, Linux provides thousands of useful and extremely powerful command-line applications. Thus, graphical user interfaces (GUIs) such as GNOME, KDE, and X Window System window managers offer applications that display a command prompt in a separate window, giving you the best of both worlds. You can be running OpenOffice Writer or GIMP in separate windows, making the most of their graphical capabilities, while you can be simultaneously displaying a command prompt in another and executing command-line utilities there.

This chapter explains the basic ideas behind command-line utilities and the commands that execute them, discusses the primary Linux command-line interpreters, explains how to access them, and explores some of the most popular utilities for running command-line utilities on your Ubuntu Linux system.

IN THIS CHAPTER

What is a shell?

Why use the command line?

Working with files and directories

Understanding Linux permissions

Getting help for Linux commands

How to configure the bash shell

NOTE **Most of the command-line examples given throughout this chapter begin with the character `$`, which is the default prompt used by Ubuntu's default Linux command interpreter. (A prompt is the character or characters that a command interpreter displays to signify that it is waiting for you to type something.) If you're following along, you should not type this character—you'll already see it (or some more complex prompt) at the beginning of your command line.**

Why Use the Command Line?

Linux provides thousands of command-line utilities that range from simple programs for creating, examining, and modifying files and file permissions to complex utilities that enable you to fine-tune the performance of your hardware and low-level operating system capabilities such as filesystems and networking. Many of these applications have graphical equivalents, especially user-level applications and system configuration utilities. For example, the System ➪ Administration menu's Networking menu item starts the graphical `network-admin` application, which performs the same functions as the command-line `ifconfig` utility. Similarly, GNOME's Nautilus file manager is roughly equivalent to a Linux command-line interpreter in general because it enables you to examine and manipulate files and directories, execute other applications, and so on.

Using the command line isn't really an alien notion even to users of operating systems such as Microsoft Windows and Mac OS X (though it's newer to the latter). Many Windows users, especially system administrators, have always found a certain convenience in some of the command-line configuration utilities provided with Windows systems, especially those related to network configuration and status reporting such as `ipconfig`. The idea of a command line may be somewhat new to long-time Macintosh users because the command line was a serious alternative to the GUI when Apple introduced its UNIX-based Mac OS X.

Even if you don't specifically want to use the Linux command line, there is one case in which you may have to, which is if your Ubuntu system encounters major hardware or configuration problems during the boot process. If, for example, your system finds that its root filesystem is corrupted or inaccessible during the boot process, you'll see a command-line prompt faster than you can say, "Hmmm, I wonder if that's a major problem." All Linux systems fall back to the command line when major problems are identified during the boot sequence to give you access to the command-line applications that you'll need to use in order to correct the problem and re-enable the system to boot normally. The Linux boot process is a command-line process, executing the appropriate configuration utilities in the correct order, the last of which is the command that starts the X Window system and your GNOME desktop (or window manager, if you've customized things). Even if your computer's filesystems are in good shape and all of your hardware and associated system software runs correctly during the boot process, configuration problems with the X Window system are a common cause of seeing a command-line prompt instead of the glorious Ubuntu login screen when you boot your system. The Ubuntu Linux boot process is described in detail in Chapter 22, "Understanding the Ubuntu Startup and Shutdown Processes."

Executing Commands from the Command Line

Applications that are designed to be executed from a Linux command line are typically referred to as command-line utilities. All command-line utilities have the same general organization—they begin with the name of the command that you want to execute, and are optionally followed by information about the way in which you want the command to behave. Anything following the name of the command that you

want to execute is known as an *argument* to that command. The arguments to each command-line utility differ based on the command that you want to execute. In an interesting example of recursion, the arguments to command-line utilities are generally referred to as *options*, each of which may take an argument.

Confused? If this is all new to you, don't worry. A few examples will help clear this up, and soon you'll be as comfortable at the command line as anyone. Let's use the Linux `ls` command as an example because it has more options than almost any other command and was also introduced in the section of Chapter 4 entitled "Understanding Linux Permissions."

The `ls` command lists information about files and directories on your Ubuntu system. The Linux `ls` command is an updated version of a classic UNIX utility by the same name. In true UNIX fashion, no one was willing to type extra characters such as `i` and `t`, so the command was given the faster-to-type abbreviation of `ls`.

When used by itself on the command line, the `ls` command simply displays the contents of the current directory, as in the following example:

```
$ ls
boot_services.txt hello.c   hello.o  include_example.c    test
hello             hello.foo hello.s  include_example.out
```

You can also supply the name of a specific file or directory as an argument to the `ls` command, as in the following examples:

```
$ ls hello
hello
```

Using the `ls` command to list the name of a file that you already know is spectacularly uninteresting (although it can be very useful when combined with wildcards, which are discussed later in this chapter). However, listing a directory shows the contents of that directory, as in the following example:

```
$ ls test
libxml2  netdev  system
```

The output from this command shows that the `test` directory itself contains three other files or directories. I happened to know that `test` was a directory—if you're not sure what types of things are in the current directory, you can use the `ls` command's `-F` option to give you this information. For example, here's the current directory as shown using the `ls -F` command:

```
$ ls -F
boot_services.txt hello.c   hello.o  include_example.c   test/
hello*            hello.foo hello.s  include_example.out
```

The `ls` command's `-F` option decorates the names of the objects in the current directory with an extra character to identify any object that isn't simply a text file. An asterisk following the name of an object shows that this is an executable file, while a slash (`/`) following the name of an object shows that this is indeed a directory.

The `ls` command's `-F` option is very useful, but (in true Linux fashion) it isn't the only way to get detailed information about each of the objects in the current directory. You can also get this sort of information using other options to the `ls` command. For example, one of the most commonly used options to

the `ls` command is the `-l` option, which means "display output in long format." Using this option gives a variety of additional information about the objects in the current directory, as in the following example:

```
$ ls -l
total 44
-rw-r--r--   1 wvh users  783 2006.03-15 06:36 boot_services.txt
-rwxr-xr-x   1 wvh users 9249 2006.03-15 06:37 hello
-rw-r--r--   1 wvh users   60 2006.03-15 06:37 hello.c
-rw-r--r--   1 wvh users   60 2006.03-15 06:37 hello.foo
-rw-r--r--   1 wvh users 2504 2006.03-15 06:37 hello.o
-rw-r--r--   1 wvh users  857 2006.03-15 06:37 hello.s
-rw-r--r--   1 wvh users  202 2006.03-15 06:37 include_example.c
-rw-r--r--   1 wvh users  736 2006.03-15 06:37 include_example.out
drwxr-xr-x   5 wvh users 4096 2006.03-15 06:37 test
```

As you can see, the long option displays more complete information about the files and directories in the current directory. From left to right, this information consists of the following: current permissions, the number of hard links to that file in the Linux filesystem (more about that later in this section), the owner and group, size, the date and time at which it was last modified, and the file or directory name.

As mentioned previously, you can combine options and arguments on the same command line to refine the behavior of most command-line utilities. For example, to get a long listing of the contents of the `test` directory, you would execute the following command:

```
$ ls -l test
total 0
drwxr-xr-x  2 wvh users 72 2006.03-15 06:37 libxml2
drwxr-xr-x  2 wvh users 80 2006.03-15 06:37 netdev
drwxr-xr-x  4 wvh users 96 2006.03-15 06:37 system
```

One other very popular option to the `ls` command is the `-a` option, which shows all of the objects in the current directory. By default, the `ls` command doesn't show objects whose names begin with a period (aka *full stop*) because all Linux directories contain two special entries that many people don't care about, but which are useful to traverse and support the hierarchical structure of a Linux filesystem. These are the "`.`" entry, which always refers to the current directory, and the "`..`" entry, which always refers to the parent of the current directory. Using the `ls -a` command to look at the contents of the current directory shows the following:

```
$ ls -a
.                  hello.c    include_example.c    .run_me_now
..                 hello.foo  include_example.out
boot_services.txt  hello.o    test
hello              hello.s    .my_music_directory
```

You'll note that the "`.`" and "`..`" entries are listed in the first column. However, you'll also note that two new files have appeared in the directory. These are the files `.my_music_directory` and `.run_me_now`, which are listed in the directory listing based on the first alphanumeric character in their names because the `ls` command ignores leading periods when sorting filenames (unless, of course, the filename has no other characters as is the case with the "`.`" and "`..`" entries, which therefore appear first in the listing).

TIP **Because files beginning with periods aren't included in directory listings unless you use the `-a` option, using filenames that begin with a period has become a standard convention for "hiding" those files. Files and directories whose names begin with a period are most commonly used to hold configuration information used by various Linux commands. The most common example of this type of file is the `.bashrc` configuration file used by the bash shell, which I'll discuss later in this chapter.**

You can combine multiple single-letter options after a single leading dash, and the `ls` command will perform all of the specified actions. For example, let's combine the `-a` and `-F` options to look at the current directory:

```
$ ls -aF
./                  hello.c     include_example.c     .run_me_now*
../                 hello.foo   include_example.out
boot_services.txt   hello.o     test/
hello*              hello.s     .my_music_directory@
```

From the output of this command, you can see that the "`.`" and "`..`" entries are both directories (actually, they are hard links to the current and parent directories), and the mysterious `.run_me_now` file is actually an executable command. The name of the `.my_music_directory` file is followed by an at symbol (@). What's up with that?

I've used the term "links" in passing previously, and this is a good time to explain the usage. Links are essentially just pointers to other files and directories in the Linux filesystem (much like shortcuts on Microsoft Windows systems or aliases on Mac OS X systems). Linux supports two types of links: hard links, which are actual connections to an existing file or directory, and symbolic links, which contain the name of some other file or directory that you want to refer to. Because hard links are actual connections to an existing data structure, the thing that you're linking to must actually exist and must be within the same disk partition as the thing that you're hard linking to. (This is a side effect of the internal data structures used by filesystems, which are too detailed to discuss here. This whole options concept may already be making you drowsy, let alone plunging into a discussion of filesystem data structures.) Unlike hard links, because symbolic links just contain the name of something else, they can point to any file or directory on your machine. As you may now suspect, an @ symbol following the name of a file or directory indicates that this file or directory is just a symbolic link to another file or directory somewhere else on your system. Let's combine the `ls` command's `-a` and `-l` options to get some detailed information about everything in the current directory:

```
$ ls -al
total 45
drwxr-xr-x   3 wvh users  400 2006.03-16 06:20 .
drwxr-xr-x  13 wvh users 1376 2006.03-16 05:38 ..
-rw-r--r--   1 wvh users  783 2006.03-15 06:36 boot_services.txt
-rwxr-xr-x   1 wvh users 9249 2006.03-15 06:37 hello
-rw-r--r--   1 wvh users   60 2006.03-15 06:37 hello.c
-rw-r--r--   1 wvh users   60 2006.03-15 06:37 hello.foo
-rw-r--r--   1 wvh users 2504 2006.03-15 06:37 hello.o
-rw-r--r--   1 wvh users  857 2006.03-15 06:37 hello.s
-rw-r--r--   1 wvh users  202 2006.03-15 06:37 include_example.c
-rw-r--r--   1 wvh users  736 2006.03-15 06:37 include_example.out
drwxr-xr-x   5 wvh users  120 2006.03-15 06:37 test
lrwxrwxrwx   1 wvh users   11 2006.03-16 06:20 .my_music -> \
    /opt2/music
-rwxr-xr-x   1 wvh users  269 2006.03-16 05:47 .run_me_now
```

This actually shows you a fair amount of information about everything in this directory, especially in terms of links. Remember that the second column in a long directory listing identifies the number of hard links to each object. For example, the second column for the "`.`" entry shows that there are three hard links to the current directory—these are the "`.`" itself, its entry in the parent directory, and the hard link to "`.`" that is present in the `test` subdirectory because the current directory is the parent of the `test` directory. The "`..`" entry, which is a hard link to the parent directory of the current directory, seems to be very popular because there are 13 hard links to it. This means that the parent directory of the current directory probably contains several other directories. Looking at the `.my_music` file, you can see that it is indeed a symbolic link because its name in the long listing actually shows the other object that it is a symbolic link to. In this case, the file `.my_music` is a symbolic link to a directory that happens to live on another filesystem that is mounted on the `/opt2/music` directory on my system.

All of the command-line options I've discussed up to this point have begun with a single dash. This isn't always necessarily the case for all commands and their options. Some Linux commands with their conceptual roots in ancient versions of UNIX (such as the `tar` command) don't even require a leading dash before a single command-line option or a single group of command-line options. This antique command-line convention is deprecated, which means that anyone who implements a command nowadays that doesn't require at least one dash before its options will be mocked and verbally abused via e-mail by the Linux and open source communities. (All UNIX commands that don't require a dash before their options have also been updated so that they can also handle finding a dash before their options.) Nowadays, command-line options always begin with a dash, but in an interesting usability twist, they can also begin with two dashes. The conventions for this are that single-letter options are preceded by a single dash, while multi-letter, "whole word" options begin with two dashes. This is necessary for two reasons:

- Most command-line utilities support both styles of options: the traditional single letter options and the newer whole-word options pioneered by the folks at the Free Software Foundation.
- Because you can combine multiple, single-letter options into a single group of options (as you've seen throughout this section), unless you use two dashes to identify a whole-word option, the command that you're executing can't differentiate between the two.

To illustrate this, let's ask for help from the `ls` command:

```
$ ls -help
/bin/ls: invalid option -- e
Try '/bin/ls  --help' for more information.
```

Well, that was actually both illustrative and friendly. When preceded with a single dash, the `ls` command interprets each of the letters that follow as a single option, and therefore complains because no option matches the letter `e` in the `-help` option. Let's try that again, correctly. Running the `ls  —help` command displays the following output (truncated here because this is just an example):

```
$ ls  --help
Usage: /bin/ls [OPTION]... [FILE]...
List information about the FILEs (the current directory by
default). Sort entries alphabetically if none of -cftuSUX
nor  -- sort.

Mandatory arguments to long options are mandatory for short
options too.
  -a,  -- all              do not ignore entries starting with .
  -A,  -- almost-all       do not list implied . and ..
[remaining output deleted]
```

Getting Information About Commands

As you can see from the final example in this section, many command-line utilities provide a `--help` option, which displays what is known as a usage message. A usage message is a short summary of how to use a command, summarizing available options and their meanings. Unfortunately, not all commands offer a `--help` option and those that do can't display all possible information about those options. As discussed in Chapters 5 and 6, the graphical user environments used on Ubuntu systems provide a huge assortment of online help that makes it easy for you to get information about how to use graphical commands. Luckily, the Ubuntu command-line environment also provides a similar amount of online help in the form of the `man` and `info` commands.

The `man` command is your best friend when using or simply experimenting with command-line programs. The `man` command displays online manual pages, formatted for your screen. In true Linux/UNIX fashion, the `man` command paginates its output using the Linux version of the familiar UNIX `more` command—known as `less`—to make it easy to scroll forward or backward in its output. Because the online reference information displayed by the `man` command is displayed a screen/page at a time and corresponds to the documentation you'd traditionally find in a printed reference manual, the documentation displayed by `man` is generally referred to as a *man page* or as *man pages*.

As an example of using the `man` command, you can type `man man` for additional information on the `man` command itself. The first part of this man page looks like the following:

```
man(1)                  Manual pager utils                  man(1)

NAME
       man - an interface to the on-line reference manuals

 SYNOPSIS
       man  [-c|-w|-tZHT  device]  [-adhu7V]  [-i|-I]
       [-m system[,...]] [-L locale] [-p string] [-M path]
       [-P pager] [-r prompt] [-S  list]  [-e extension]
       [[section] page ...] ...
       man  -l  [-7] [-tZHT device] [-p string] [-P pager]
       [-r prompt] file ...
       man -k [apropos options] regexp ...
       man -f [whatis options] page ...

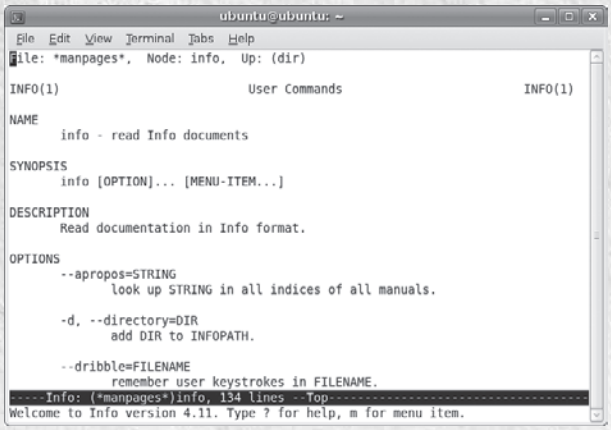
DESCRIPTION
       man is the system's manual pager. Each page argument given to
       man is normally  the  name  of  a  program,  utility or func-
       tion. The manual page associated with each of these arguments
       is  then found  and  displayed. A section, if  provided, will
       direct  man  to look only in that section of the manual.  The
       default action is to search  in  all  of the  available  sec-
       tions,  following  a  predefined order  and to  show only the
       first page found, even if page exists in several  sections.
[Additional output removed]
```

continued

continued

As you can see, the online man pages for Linux command provide extensive information about available options, but also provide a significant amount of general information about using the commands themselves, usually including examples.

A more modern alternative to the `man` command, the `info` command provides similar usage and explanatory information about Linux commands. The `info` command uses Emacs to display its output and also provides more extensive and up-to-date command documentation than the traditional `man` command. The following figure shows a window displaying info on the `info` command.



The info command running in a terminal window

The GNU folks introduced the `info` command to provide a richer, more robust environment for displaying online help for command-line utilities. The text files that are used by the `info` command provide a very rich syntax for hyperlinking between different sections of info documentation, and the use of Emacs as the application for displaying info pages provides a much more flexible environment for moving around in the info that you're looking at. The `info` command can also display traditional Linux `man` pages, so you can even examine existing man pages using `info`—the `info man` command is perhaps the best example of that sort of recursion.

The Linux `man` and `info` commands will quickly become your new best friends when you are using or simply learning about command-line utilities.

What's a Shell?

A shell is the generic name given to any Linux command-line interpreter, and comes from UNIX, the conceptual parent of Linux. UNIX was one of the first operating systems to introduce the idea of using a command-line interpreter that was not built into the operating system and which had no special permissions to do mysterious operating system tasks. These ideas have been preserved in every conceptual descendant of UNIX and have proved handy for several reasons. The most interesting of these is that,

because it is a separate, standalone executable, a Linux system can offer each user their choice of multiple shells, all of which can be upgraded independently from the operating system.

The original UNIX shell was written by Ken Thompson, one of the two primary creators of UNIX, and was extended by John Mashey, also at Bell Labs. These shells were somewhat primitive—the first shell that has the types of capabilities we have come to know and love today was the Bourne Shell, written by Stephen Bourne at Bell Labs in 1974. This shell, known to its friends as `/bin/sh`, is the shell that is used by default on most UNIX systems, and is the conceptual parent of the `/bin/bash` shell (which stands for Bourne-Again Shell) that is used by default on most Linux systems today. As I'll explain in the next section, there are a variety of other shells to choose from on Linux systems. If you've come to Linux from a version of UNIX that features another shell, you'll be able to feel at home pretty quickly (and perhaps even bring over and reuse your existing shell configuration files).

TIP **A term that you will frequently hear when discussing shells, Linux, or UNIX in general is the term *shell script*. A shell script is a command file containing commands that are either internal to the shell (such as those for setting variables, conditional expressions, and looping constructs) or which reference other commands on your Linux system. The ability to write very sophisticated command files is a feature of the Linux/UNIX utility model, where each utility does one thing and does it well, and many different utilities can be linked together via pipes or temporary files to process each other's output.**

For additional information about shells on UNIX systems, see:

- `www.softpanorama.org/People/Shell_giants/introduction.shtml`
- `www.faqs.org/faqs/unix-faq/shell/shell-differences/`
- `www.unix.org.ua/orelly/unix/ksh/ch01_03.htm`

Available Shells for Linux Systems

A standard Ubuntu distribution installs only the standard GNU bash shell. However, other shells are available for installation through `apt-get` or the Synaptic Package Manager. The following shells are found on many Linux systems:

- `/bin/ash`: The Almquist shell, a small-footprint shell that began as a clone of the SYSV R4 version of the Bourne shell. The ash shell is often used on embedded systems or during the system startup and installation process on some Linux distributions.
- `/bin/bash`: The default, Bourne-Again shell inspired by `/bin/sh` and `/bin/ksh`. This is the default shell installed and used by all users on Ubuntu systems.
- `/bin/csh`: If you install tcsh (see its description later in this list), a symbolic link from `/bin/csh` to `/bin/tcsh` is also created. This link is provided to support shell scripts that reference the traditional location of the C-shell, the standard shell used on Berkeley UNIX (BSD) systems and their descendants. The C-shell supports configuration commands that are reminiscent of C-language programming constructs, and was originally written by Bill Joy, later a founder of Sun Microsystems.
- `/bin/dash`: The Debian Almquist shell, the Debian Linux distribution's version of the standard Almquist shell. If you're interested in experimenting with this shell, it is available in the standard Ubuntu repository.

continued

continued

- `/bin/ksh`: If you install `pdksh` (see its description later in this list), a symbolic link from `/bin/ksh` to `/bin/pdksh` is also created. This link is provided to support shell scripts that reference the traditional location of the Korn shell, the standard shell used on later SYSV UNIX systems from AT&T and their descendants.
- `/bin/nash`: Another small-footprint shell used during the startup process on many Linux systems, specifically on Red Hat and Fedora Core Linux systems.
- `/bin/pdksh`: An open source version of David Korn's Korn shell, written at AT&T and was available with the SYSV R3 (as part of the "Experimental Toolchest") and R4 (as a completely supported utility) UNIX distributions. The Korn shell (and thus `pdksh`) is completely backward-compatible with the original Bourne shell. If you're interested in experimenting with this shell, it is available in the standard Ubuntu repository.
- `/bin/sh`: A symbolic link to `/bin/bash`, provided for compatibility with generic Linux and UNIX shell scripts.
- `/bin/tcsh`: The TENEX C-shell, which is an advanced, open source version of the C-shell with command-line editing extensions that were originally introduced by the command interpreter used by DEC's TOPS-20 operating system for PDP-10 systems, which began life as BBN's TENEX (for the 10 in PDP-10) operating system and was therefore later mutated into TWENEX (for the 20 in TOPS-20) by the PDP-10 hacker community. If you're interested in experimenting with this shell, it is available in the standard Ubuntu repository.
- `/bin/zsh`: The Z shell is a powerful, tremendously extensible shell that provides many of the capabilities of bash and ksh, and much, much more. If you're interested in experimenting with this shell, it is available in the standard Ubuntu repository.

You can use the `chsh` command-line command to change the default shell used by your account if you decide that you want to use a shell other than the one that is currently listed in your `/etc/passwd` entry. However, you can't change your login shell to any random binary—all of the programs to which you can change your shell are listed in the text file `/etc/shells` on your Ubuntu system. If you want to be able to set some other application as a login shell using the `chsh` command, you must first add it to the file `/etc/shells`.

Getting to a Shell

Assuming that I've hyped the value of the Linux command line sufficiently and that your Ubuntu system boots correctly in graphical mode, you're probably wondering how to start a shell so that you can experiment a bit. Ubuntu provides two applications that are the most common mechanisms for starting a shell, one that is a standard part of the X Window system distribution, and another that is a standard part of any GNOME distribution. Both of these are referred to as terminal applications, not because they imply an end to life as we know it but because they are reminiscent of the user experience on systems without

bitmapped graphical displays, when most users accessed their computer systems through terminals that couldn't do much more than display the input and output of a command interpreter.

Using the GNOME Terminal Application

The GNOME terminal application is the most common way of starting a command-line shell on a graphical Ubuntu system. To start the GNOME Terminal, select the Terminal command from the Applications ➪ Accessories menu. Figure 14-1 shows the default GNOME Terminal window.

FIGURE 14-1

The GNOME Terminal application



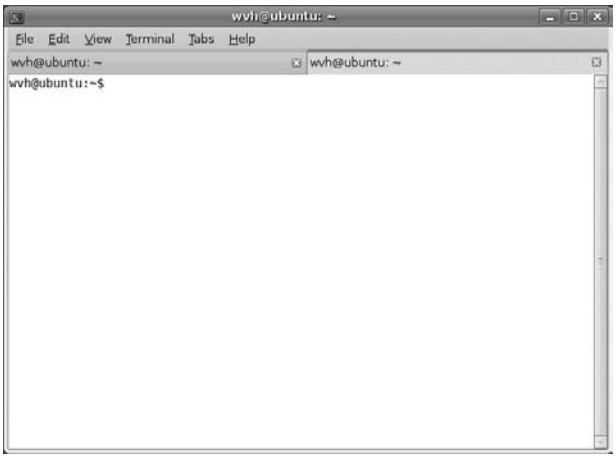
As with most GNOME applications, the GNOME Terminal provides extensive online help, which is available by selecting Help ➪ Contents.

The GNOME Terminal application provides a variety of menus that make it easy to configure things such as the title of a GNOME Terminal window (Terminal ➪ Set Title); configure the character set (Terminal ➪ Set Character Encoding); and configure the size, fonts, colors, and other display attributes of the GNOME Terminal window (Edit ➪ Current Profile). In my opinion, its most generally useful feature is its ability to open and manage multiple command-line sessions within a single Terminal window. To open a new tab, select the File ➪ Open Tab command. A new tab displays, as shown in Figure 14-2.

Once multiple tabs are displayed, you can move between tabs by clicking on a specific tab or by using the appropriate command from the Tabs menu. If you want to have an existing tab display in its own window, you can select the Tabs ➪ Detach Tab command. You can also start a new GNOME Terminal window by selecting the File ➪ Open Terminal command.

FIGURE 14-2

Multiple tabs in the GNOME Terminal



Using the KDE Konsole

The terminal application provided on Kubuntu systems is known as Konsole, both because its command-line behavior mimics the behavior of the console on an older UNIX or Linux system, and also because the names of all good KDE applications begin with the letter K.. The Konsole application is the most common way of starting a command-line shell on a graphical Kubuntu system. To start the Konsole, select the Konsole command from the K Menu ➪ System menu on a KDE 3-based system, or the Terminal command from the K Menu ➪ Applications ➪ System menu on a KDE 4-based system. Figure 14-3 shows the default Konsole window on a KDE 3-based Kubuntu system. Konsole looks almost exactly the same on a KDE 4-based Kubuntu system, modulo the window decorations.

As with most KDE applications, the Konsole provides extensive online help as a handbook that is available by selecting Help ➪ Konsole Handbook.

You can configure things such as the character set, font size, fonts, colors, and other display attributes of the Konsole application using the options on the Settings menu. You can also configure many other options by selecting the Settings ➪ Configure menu command to display the dialog shown in Figure 14-4.

In my opinion, its most generally useful feature is its ability to open and manage multiple command-line sessions within a single Terminal window. To open a new tab, select the Session ➪ New Shell command. A new tab displays, as shown in Figure 14-5.

Once multiple tabs are displayed, you can move between tabs by clicking on a specific tab or by using the appropriate tab from the bottom of the View menu. If you want to have an existing tab display in its own window, you can select the View ➪ Detach Session command. You can also start a new Konsole window by selecting the Session ➪ New Window command.

FIGURE 14-3

The KDE Konsole application

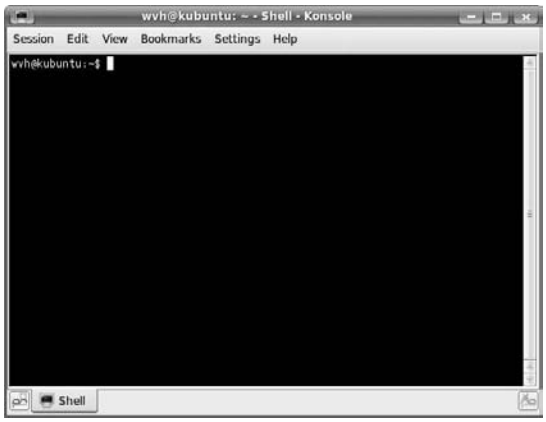


FIGURE 14-4

Configuring the KDE Konsole

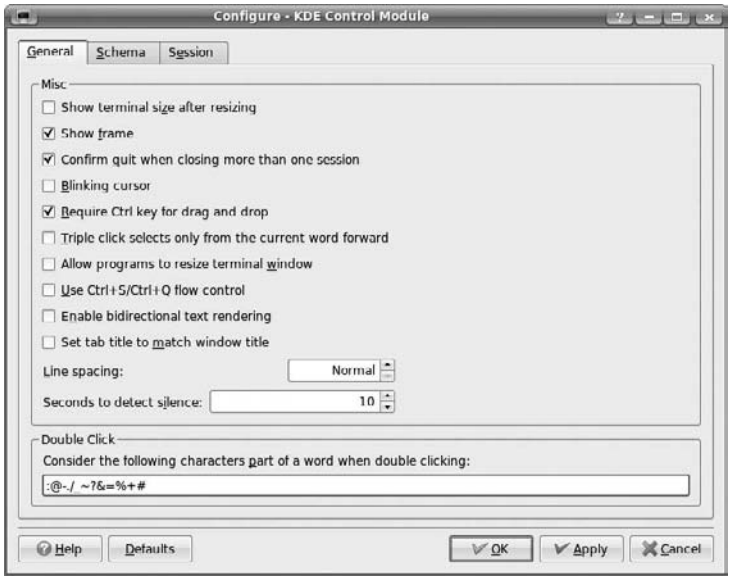
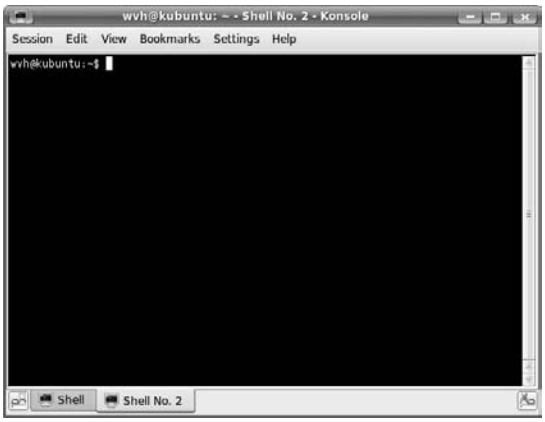


FIGURE 14-5

Multiple tabs in the KDE Konsole

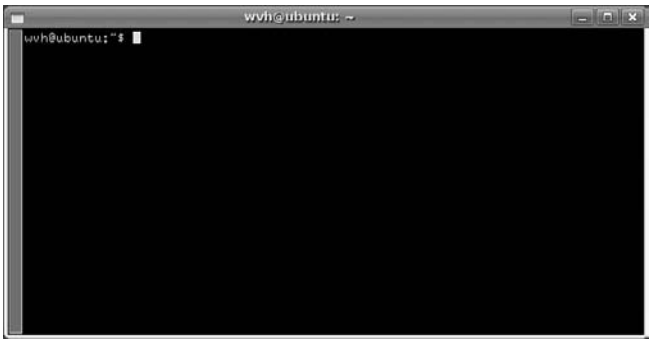


Using the X Window System Terminal Application

Any X Window System distribution includes a lightweight terminal application known as `xterm`. The executable for the X Window System's terminal application, shown in Figure 14-6, is located in `/usr/bin/xterm`.

FIGURE 14-6

A traditional xterm



To start the `xterm` application, do any of the following:

- Add a menu item for the `xterm` application, as explained in "Customizing GNOME Menus" for GNOME in Chapter 5, and for KDE in Chapters 7 and 8. You can then select that menu item to run the command.

- Create a launcher for the `xterm` application on your panel using the Run Application panel applet, and then select that panel applet. Adding items to a GNOME panel is explained in "Customizing Panel Contents" in Chapter 5. Adding items to a KDE 3 panel is explained in "Customizing Panel Contents" in Chapter 7. Adding items to a KDE 4 panel is explained in "Using the KDE 4 Panel" in Chapter 8.
- Create a desktop launcher for the `xterm` application by right-clicking on your desktop's background, selecting the Create Launcher command, and filling out the dialog shown in Figure 14-7. You can then start an `xterm` by clicking on the desktop icon for the launcher.
- Execute the `xterm` command directly from the GNOME Terminal, which is somewhat recursive, but eminently doable.

FIGURE 14-7

Creating a desktop launcher for the xterm application



Many users who are new to Ubuntu but are already familiar with the `xterm` application find that they prefer familiar `xterm`s to the GNOME Terminal. For complete details on all of the capabilities of the `xterm` application, see its man page—perhaps by starting an `xterm` and executing the `man xterm` command in that window.

As a generic X Window system application rather than a GNOME or KDE application, you'll note in Figure 14-3 that an `xterm` window doesn't provide a menu or toolbar to provide easy access to configuration commands. Instead, you display `xterm`'s configuration menus by holding down the Ctrl key on your keyboard and clicking the left, middle, or right mouse buttons on the background of the `xterm` window. Figure 14-8 shows the menu that you see when holding down the Ctrl key and left-clicking the background of an `xterm window`.

Each `xterm` configuration menu provides access to different `xterm` configuration commands:

- **Ctrl+left-click:** The Main Options menu, which primarily enables you to send various signals and special key sequences to the selected `xterm`.
- **Ctrl+middle-click:** The VT Options enables you to set various characteristics of the terminal emulation provided by the selected `xterm`. The name VT Options comes from the name of the default type of terminal that is emulated by an `xterm`, which is a DEC VT-100. Figure 14-9 shows the VT Options menu.
- **Ctrl+right-click:** The Fonts menu enables you to control the size of the fonts used in the selected `xterm`. Figure 14-10 shows the Fonts menu.

FIGURE 14-8

An xterm's Main Options menu

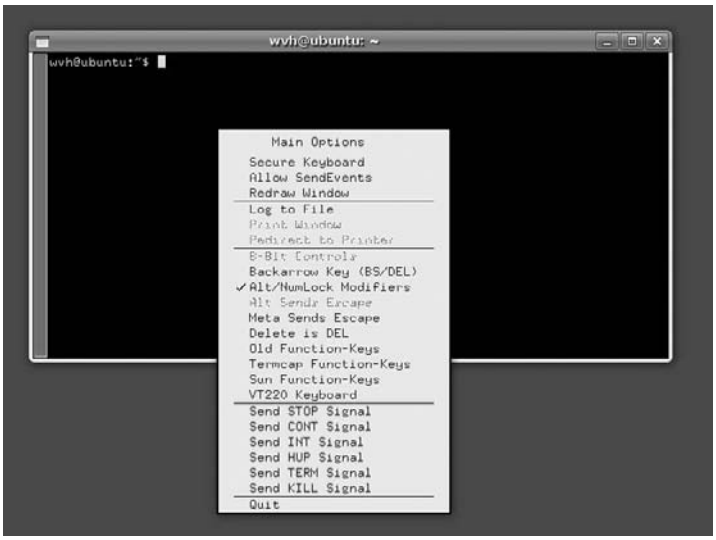
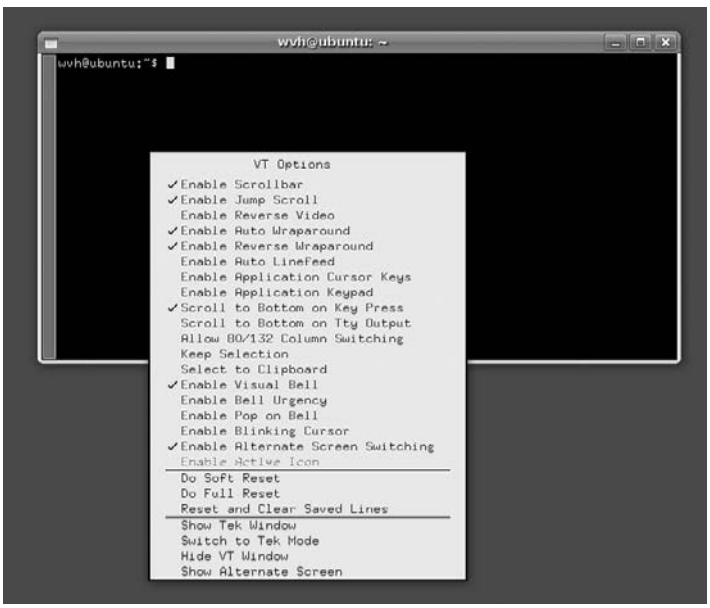


FIGURE 14-9

An xterm's VT Options menu



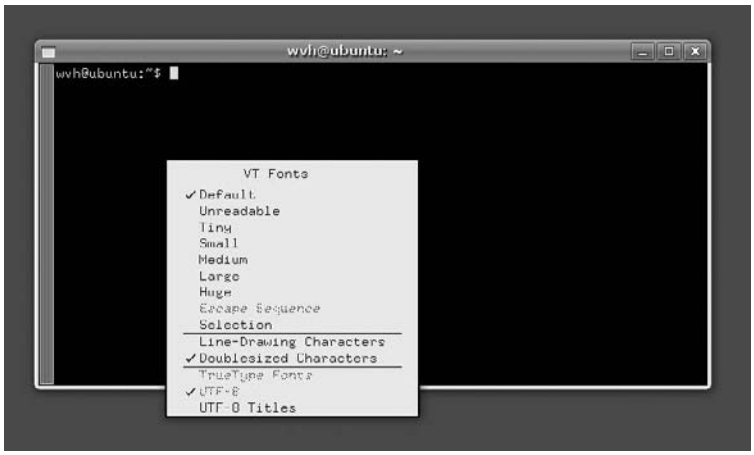
TIP **By default, an `xterm` starts a copy of your login shell, but one of the most common ways of using the `xterm` application is to automatically start another application in the xterm window using the `xterm` command's `-e` (execute) option and specifying the name of the application that you want to start. For example, the following command starts an `xterm` that is running the `ssh` command to connect to the system `writing.vonhagen.org`:**

```
$ xterm -e ssh writing.vonhagen.org
```

This command would prompt for a login password, and then give me a shell on the specified system. (For more information about connecting to remote systems, see Bonus Chapter 1, "Connecting to Other Systems.") Because xterms don't feature multiple tabs like the GNOME Terminal application (explained in the previous section), an `xterm -e` command is one of the most common ways for Ubuntu users to start a separate window that is running another nongraphical command.

FIGURE 14-10

An xterm's Fonts menu



Popular Command-Line Commands

The section of this chapter that introduced command-line utilities, options, and arguments used the standard Linux `ls` command as an example. Listing files and directories is one of the most common things that you'll want to do from a shell (or, at least, will find yourself doing). However, as explained in the previous section, your Ubuntu system provides hundreds of commands. I can't go into all of them here because that would mean that Wiley would need to put wheels on this book, but there are certainly some common tasks that you'll want to do from a shell. This section highlights some of the more common tasks you'll want to do from the command line, and concludes by discussing how to figure out which commands to use if you don't already know the cryptic name of the appropriate command-line utility.

Moving Around in the Linux Filesystem

Just as when you use a graphical user interface, you will want to organize your files and directories in some logical fashion if you're working from the command line. Once you've organized things that way, you'll need to use the `cd` (change directory) command, which does exactly what you'd expect. In hard-core Linux terms, the `cd` command changes your current working directory to be the directory that you supply as an argument. For example, to change directory into a directory named `test`, you'd simply execute the following command:

```
$ cd test
```

TIP **If you're hanging around with hard-core Linux folks and want to fit in, refer to changing to another directory as *cd'ing* there.**

If you don't know the name of the directory where you are currently working, that's somewhat scary, but you can use the `pwd` command (process working directory) to provide that information. The `pwd` command lists the full pathname of the directory in which the current process (in this case, your shell) is running, as in the following example:

```
$ pwd
/home/wvh/test
```

You can change to any directory on your Linux system by specifying its full pathname as the argument to the `cd` command, as in the following example:

```
$ cd /usr/local
$ pwd
/usr/local
```

As you can see, my shell is now working in the `/usr/local` directory and any other commands that I execute will be performed in that directory. As a convenience, if you execute the `cd` command without any arguments, the `cd` command automatically changes directory to your home directory, which is probably where most of your files and directories are located in the first place.

TIP **As you work at the command line, you may find that you often want to change to some other directory, do one or two things, and then return to the directory that you started in. When changing directories frequently, it's handy to let your shell keep track of things like "where did I come from." To save you time and mental energy, most Linux shells provide built-in commands called `pushd` and `popd` to track this for you. The `pushd` command changes to a specified directory, but remembers where you came from by pushing the name of the directory that you came from onto a list of previous locations known as a *stack*. Once you're done working there, you can type the `popd` command to return to the directory that you came from, which removes it from the stack (i.e., the list of previous working directory locations). If you're unfamiliar with the idea of a stack, the standard example of a stack is the spring-loaded plate stacks of plates that you find at a cafeteria or your favorite local buffet. Pushing something onto the stack is analogous to putting a new plate on the top of the stack of available plates. Popping something off the stack is analogous to removing the top plate from the stack, which makes the next one available.**

Copying, Moving, Renaming, and Deleting Files and Directories

Copying, moving, and renaming files are some of the most common actions that any computer user performs, regardless of whether you're using the command line or working with a graphical user interface.

Whether you're cloning an existing file to use the former as a starting point for new work or reorganizing your files and directories to make things easier to find, copying and moving files and directories is just as easy to do from the command line as it is from a GUI, modulo the fact that you have to type rather than point and click. In many cases, the zillions of options provided by many Linux command-line utilities give you much finer control over things than a generic graphical command ever could.

In continuing the Linux/UNIX tradition of "why type more letters than necessary," the Linux command to copy files and directories is the `cp` command. To copy one file to another, you simply execute the `cp` command, followed by the name of the existing file and the name that you want to have the copy to have. For example, to copy a file named `file.txt` to the new file `new_file.txt`, you would execute the command:

```
$ cp file.txt new_file.txt
```

CAUTION **Be very careful when specifying the name of the file that is the destination of a `cp` command. If a file with the same name as the file to which you are copying already exists, the `cp` command overwrites the existing file with the contents of the file that you are copying. If you're really paranoid about this, see the section later in this chapter on "Defining and Using Aliases" for information about how to permanently set up a safe version of the `cp` command.**

By specifying the full pathnames of the source or destination file, the `cp` command can copy files anywhere in the Linux filesystem in which you have permission to create files. For example, to copy the file `file.txt` from my home directory to the file `newfile.txt` in the directory `/tmp` (where everyone can create files), I could execute the following command regardless of the file or directory in which I was currently working:

```
$ cp /home/wvh/file.txt /tmp/newfile.txt
```

You can also use the `cp` command to copy a file to different directories by specifying the name of a directory as the target of the `cp` operation. Continuing with the previous example, if I had wanted to copy the file `file.txt` from my home directory to `/tmp`, but have the copy retain the same name as the original file, I could have done so with the following command:

```
$ cp /home/wvh/file.txt /tmp
```

This command would create the file `/tmp/file.txt` as an exact copy of `/home/wvh/file.txt`.

The `cp` command doesn't require any special options when simply copying files. However, you can also copy directories using the `cp` command, by specifying its `-r` (recursive) option. For example, to copy the contents of a directory of mine called `include_test` to a new directory called backup, I would execute the following command:

```
$ cp -r include_test backup
```

After the command completed, the new backup directory would be an exact copy of the `include_test` directory. If you're from Missouri and need proof that the right thing is happening, you can add the `-v` (verbose) option to the `cp` command to list each part of the copy command. This command and its output would look like the following:

```
$ cp -rv include_test backup
'include_test' -> 'backup'
'include_test/libxml2' -> 'backup/libxml2'
'include_test/libxml2/xmlops.h' -> 'backup/libxml2/xmlops.h'
'include_test/netdev' -> 'backup/netdev'
'include_test/netdev/devname.h' -> 'backup/netdev/devname.h'
```

```
'include_test/system' -> 'backup/system'
'include_test/system/libxml2' -> 'backup/system/libxml2'
'include_test/system/libxml2/xmlops.h' -> \
    'backup/system/libxml2/xmlops.h'
'include_test/system/netdev' -> 'backup/system/netdev'
'include_test/system/netdev/devname.h' -> \
     'backup/system/netdev/devname.h'
```

As you'd hope, a recursive `cp` command also recursively copies all subdirectories and any files that they contain.

I mentioned earlier that specifying the name of a directory as the target when copying a file creates a file with the same name as the original in the target directory. The same is true when specifying the name of a directory as the target for a recursive `cp` command. This can initially be confusing if you're trying to copy a directory to a directory with a new name and a directory with that name already exists. In that case, the `cp` command creates a directory with the same name as your original directory as a subdirectory of the target directory.

Moving and renaming files from the command line is similarly easy. Linux systems don't actually provide a separate rename command—on a Linux system, the `mv` command both moves existing files and directories to a new location and also optionally renames them as part of the move operation. For example, to rename the file `file.txt` to the new name `newfile.txt`, you would do the following:

```
$ mv file.txt newfile.txt
```

The `mv` command follows the same conventions as the `cp` command in terms of specifying full pathnames and when specifying a directory as the name of the location to which you want to move a file. For example, the following command moves the file `/home/wvh/file.txt` to the `/tmp` directory, and renames it `newfile.txt`:

```
$ mv /home/wvh/file.txt /tmp/newfile.txt
```

If you wanted to move the file to the `/tmp` directory and keep its original name, you would do so with the following command:

```
$ mv /home/wvh/file.txt /tmp
```

Unlike the `cp` command, the `mv` command doesn't need any special flags when moving or renaming directories. For example, to rename the directory `test` to backup, you would do the following;

```
$ mv test backup
```

As with the `cp` command, if the backup directory already existed, the `mv` command would move the `test` directory into the backup directory, retaining its original name. The `mv` command also provides a `-v` (verbose) option to show you what's going on during an `mv` operation, but this is somewhat less exciting when just renaming a directory, as you can see from the following example:

```
$ mv -v test backup
'test' -> 'backup'
```

In this case, the `mv` operation only needs to rename the directory. However, because of the way that Linux filesystems are created, moving directories across filesystems that live on different hard disk partitions

requires that each file and subdirectory first be copied and then deleted by the `mv` command, as in the following example:

```
$ mv -v include_test /tmp/backup
'include_test' -> '/tmp/backup'
'include_test/libxml2' -> '/tmp/backup/libxml2'
'include_test/libxml2/xmlops.h' -> '/tmp/backup/libxml2/xmlops.h'
'include_test/netdev' -> '/tmp/backup/netdev'
'include_test/netdev/devname.h' -> '/tmp/backup/netdev/devname.h'
'include_test/system' -> '/tmp/backup/system'
'include_test/system/libxml2' -> '/tmp/backup/system/libxml2'
'include_test/system/libxml2/xmlops.h' -> \
   '/tmp/backup/system/libxml2/xmlops.h'
'include_test/system/netdev' -> '/tmp/backup/system/netdev'
'include_test/system/netdev/devname.h' -> \
   '/tmp/backup/system/netdev/devname.h'
'include_test/new_file' -> '/tmp/backup/new_file'
removed 'include_test/libxml2/xmlops.h'
removed directory: 'include_test/libxml2'
removed 'include_test/netdev/devname.h'
removed directory: 'include_test/netdev'
removed 'include_test/system/libxml2/xmlops.h'
removed directory: 'include_test/system/libxml2'
removed 'include_test/system/netdev/devname.h'
removed directory: 'include_test/system/netdev'
removed directory: 'include_test/system'
removed 'include_test/new_file'
removed directory: 'include_test'
```

As you can see, much of this output is very similar to that of the `cp` command. In this example, I first used the `df` command to demonstrate that my current directory ("`.`") and the target directory (`/tmp`) are actually located on different filesystems. You can see the difference in the verbose output of the `mv` command.

Along with copying and moving files and directories, deleting files and directories is a similarly common operation. Whether you're deleting things to free up disk space for future projects, deleting backup copies of projects that you've finished, or deleting things that you don't want anyone else to see, it's quite easy to delete files from the command line. The key to doing so is the `rm` command, which continues the "the less typing, the better" philosophy and stands for remove.

To remove a single file on a Linux system, you simply type the `rm` command followed by the name of the file. For example, to delete the file `file.txt`, you would do the following:

```
$ rm file.txt
```

The file is gone. As with the `cp` command, the `rm` command provides a `-r` (recursive) option to enable you to remove entire directories. For example, to permanently remove a directory named `test`, you would issue the following command:

```
$ rm -r test
```

TIP **You can remove empty directories using the `rmdir` command, but I don't find that to be a common scenario. I rarely have empty directories just sitting around because I usually create directories to store something in them. I generally use the `rm -r` command to delete any directory because that way I don't have to manually delete its contents first. Your mileage may vary.**

If you've moved to Ubuntu or Kubuntu from a Microsoft Windows system, you may be painfully aware that it is easy to recover files that you've deleted on a Windows system because it initially just erases the directory entry that identifies the file or folder that you're deleting. This is not the case on a Linux system. When you delete a file or directory on a Linux system, all of the disk storage associated with the file, directory, and the contents of that directory are returned to a general list of free space that is available on your system. Although deleted files and directories can still be recovered on a Linux system, it is much harder to do so and requires the assistance of someone who really knows the details of the filesystem.

The downside of this is that you can't easily recover any files that you've deleted by mistake unless you do so using a graphical utility such as a file manager. To help protect yourself from accidentally deleting files, see the section later in this chapter on "Defining and Using Aliases" for information about how to permanently set up a safe version of the `rm` command that will ask you for confirmation before it actually removes anything.

TIP **All of the `cp`, `mv`, and `rm` commands have many more options than I've covered here. To see all of the options available for any of these commands, use the `man` or `info` commands to see the online reference information for the command that you're interested in.**

Changing File and Directory Permissions

Chapter 4 introduced the idea of Linux users and groups, and also provided an initial explanation of how file and directory permissions are represented when listing the contents of files and directories. You saw the same thing earlier in this chapter when exploring the `ls` command's `-l` option. Being able to protect files so that only certain users or groups can examine or execute them is certainly handy, but it begs the question, "How do I do that?" You can easily change file and directory permissions from graphical file managers such as GNOME's Nautilus or KDE's Konqueror or Dolphin by right-clicking, selecting Properties, and selecting the Permissions tab, as shown in Nautilus in Figure 14-11. Fortunately, changing file and directory permissions from the command line is similarly easy, and may even be easier.

The `chmod` (change mode) command enables you to change any of the permissions associated with files and directories that you own. The `chmod` command supports two ways of setting permissions:

- Using octal values that actually represent the bits in the permission field of a file or directory.
- Using friendly letters to represent permissions for the user who owns a file or directory (`u`), the group with which that file or directory is associated (`g`), other users on the system that are not in that group (`o`), or all users on the system (`a`), and using the plus and minus signs to add or subtract permissions. This is known as changing permissions in symbolic mode.

In this section, I'll focus on the second approach because many of us may be a bit rusty working in octal, and the symbolic mode is easier to visualize and use. I'll use as an example the file `hello.sh`, which happens to be a bash command file (shell script) that simply prints the string "Hello, World!" in the shell window from which it is executed. Let's first look at its default permissions:

```
$ ls -l hello.sh
-rw-r--r--   1 wvh wvh 36 2006.03-18 16:30 hello.sh
```

FIGURE 14-11

Setting permissions in Nautilus



As explained in detail in Chapter 4, the characters at the left portion of the `ls` command's long output describe the current file permissions. These file permissions fall into four groups: the leading dash identifies the type of object that this is (a dash means that it's just a regular file), and the remaining three groups of three characters indicate the permissions that the owner of the file (wvh) has, the groups associated with that file (also named wvh), and the rights that all other users have to the file. Each of these groups of three characters can have the following values (in order):

- `r`: The file or directory can be read.
- `w`: The file or directory can be written to.
- `x`: The file can be executed or the directory can be searched for other files or directories.
- `-`: The permission is not granted.

The two instances of "wvh" in the `ls` output, from left to right, identify the owner of the file and the group that is associated with the file. In the sample `ls` output, we therefore can see that I (as the owner of the file) can both read and write to the file, but that other members of the group wvh can only read the file, and that all other users on the system can also only read the file.

Because this is a shell script, I presumably want to be able to execute it. To make the file executable by its owner (me), I would issue the following command:

```
$ chmod u+x hello.sh
```

Listing the file again, we can see that the file is now marked as being executable by its owner:

```
$ chmod u+x hello.sh
$ ls -al hello.sh
-rwxr--r--   1 wvh wvh 36 2006.03-18 16:30 hello.sh
```

If I want others in the group wvh and all other users on the system to be able to execute this shell script, I could issue the following commands to make those changes and verify them:

```
$ chmod g+x hello.sh
$ ls -al hello.sh
-rwxr-xr--   1 wvh users 36 2006.03-18 16:30 hello.sh
$ chmod a+x hello.sh
$ ls -al hello.sh
-rwxr-xr-x  1 wvh users 36 2006.03-18 16:30 hello.sh
```

That's a lot of typing for a repetitive series of operations! Luckily, the `chmod` command enables you to group multiple permission operations. Suppose I wanted to remove the ability for anyone other than myself to execute this script. I could do this with the following single command:

```
$ chmod go-x hello.sh
$ ls -l hello.sh
-rwxr--r--   1 wvh users 36 2006.03-18 16:30 hello.sh
```

Note that I used `o`, which means anyone who is not in the group wvh, rather than `a`, which means all users on the system, because all users on the system would include me.

You can add or remove read and write permissions for any supported permission group using the `chmod` command and similar options. Files and permissions have the same meanings for directories as they do for files, with the exception of the executable bit, which for directories grants permission to search that directory. For example, you can create a directory that no one but you can read, but still enable people to execute programs or shell scripts that it contains by removing read permission (`r`), but making sure that search permission (`x`) is still set. People won't be able to see what's in the directory, but they can execute scripts or programs that it contains if they know their exact names.

TIP **On a related subject, you can use the `chown` (change owner) and `chgrp` (change group) commands to change the owner and group of a file, but that's a privileged operation that you must do using `sudo`, and is a relatively rare task unless you're working as a system administrator. The online reference pages for these commands explain how to run them and provide a good number of examples.**

Default Permissions When Creating Files and Directories

A umask (user mask) is the classic UNIX mechanism for setting the default protections of file and directories that you create. By default, the umask is a four-digit octal number that is logically ANDed with the generic file protections of octal 0666 (ironically) when you create a file, or octal 0777 when you create a directory. The default umask value on most Linux distributions is 0002, meaning that any file you create is created with the octal protection mode 0664—both the owner and group can read and write to any file that you create, but randoms can only read the file. Similarly, any directory that you create is created with the octal protection mode 0775—both the owner and group can create files in that directory, and anyone can list the contents of the directory and search for files in it. You find out a user's default umask setting by issuing the `umask` command from any Linux command-line prompt.

Most people set their umask to 0022 in their shell configuration command file (typically `~/.bashrc`) to change their default file creation settings so that files can only be written by their owners (i.e., are created with an octal protection of 0644) and directories can only be written to by their owners (i.e., are created with an octal protection of 0755).

Finding Commands That Do Specific Things

The command line is a pretty friendly place to execute commands if you know what you're doing and know what command to do it with. Unfortunately, given the Linux tendency to avoid long command names by dropping extra characters from their names, it's often hard to figure out what command you want to use to do a specific task.

Luckily, the Linux online manual command, `man`, provides a truly convenient keyword search option (`-k`) that automatically searches for a specified keyword or phrase in the description fields of each man page. For example, suppose that you didn't know about the `chmod` command, and wanted to see a list of commands that has something to do with file permissions. You could look for appropriate commands using the following command:

```
$ man -k permissions
chmod (1)            - change file access permissions
```

Note that this can be tricky because just like a Web search, you have to look for the right phrase. However, at least the possibility is always there. You can search for single words or quoted strings to try to refine the list of commands that match what you're looking for.

If your `man -k` output lists a command but you can't seem to execute the command (you see the message "command not found" when you try to execute it), this means that the directory where that command is stored isn't one of the directories where you ordinarily look for commands. In that case, you can try using the `whereis` command to locate that command, as in the following example:

```
$ man -k "physical volume"
pvchange (8)         - change attributes of a physical volume
pvdisplay (8)        - display attributes of a physical volume
pvremove (8)         - remove a physical volume
pvs (8)              - report information about physical volumes
pvscan (8)           - scan all disks for physical volumes
vgextend (8)         - add physical volumes to a volume group

$ pvscan
bash: pvscan: command not found

$ whereis pvscan
pvscan: /sbin/pvscan /usr/share/man/man8/pvscan.8.gz
```

The `whereis` command tells you where the command is located (and where its documentation is found). You can then execute that command by specifying the full pathname of the executable, which in this case is `/sbin/pvscan`.

Working with the Bash Shell

Now that you've started a shell and learned some of the commands that you want to run from the command line, it's time to explore some of the cooler aspects of working with and configuring the bash shell. Obviously, you could simply type commands at the command prompt to your heart's content, but the bash shell provides many nice features that make it easy to use and are well worth taking advantage of. Some of these features can save a significant amount of time, typing, and typos when executing the same command over and over, when executing a group of similar commands, when typing the names of long commands, or complex filenames as arguments to commands, and so on.

Using Command History

All Linux shells provide a history mechanism, which is a way of recalling and re-executing commands that you have previously executed. To see the list of recently executed commands, you can simply type `history` at any shell prompt, which displays something like the following:

```
$ history
231  make distclean
232  ../gcc/configure  -- enable-threads
233  make bootstrap
234  info -f gcc.info
235  export INFOPATH='pwd':$INFOPATH
236  vi ~/.bashrc
237  pushd ~/new
237  ls /mnt/terastation/CDs
238  pushd /mnt/terastation/CDs
239  find . -name "T*.txt"
240  find . -name "T*.txt" -exec rm {} \;
250  rar x ../United_Empire_Loyalists.rar
[much more output deleted]
```

You'll have to trust me that this list contains a recording of a certain number of the commands that I have executed on my system. The number of commands that the history mechanism records is identified by the `HISTSIZE` environment variable. (The section "Using Environment Variables" later in this chapter provides information about setting and displaying the values of environment variables.)

You can recall or re-execute commands in several ways:

- Using the shell's built-in `fc` command
- Using the shell's mechanism for interacting with the commands that you have typed (known as the *readline library*) to scroll backward through the history list from the command line
- Using the exclamation point to identify the command that you want to display and re-execute

I generally use one of the last two mechanisms, so I'll focus on those here.

You can redisplay any recently executed command by typing the `Ctrl+P` command at any command prompt. Each time you type `Ctrl+P`, the shell displays the previous command in its history list. When you find the command that you want to re-execute, you can simply press Return to re-execute it or edit that command using Emacs commands such as `Ctrl+B` to move backward one character in that command, `Ctrl+F` to move forward in that command one character, the backspace and delete keys to delete the previous and current characters, and so on.

Obviously, if the command that you want to re-execute is far back in your command history list, scrolling back through a few hundred commands can be a bit tedious. In this case, you can use the exclamation point to identify a particular command from the history list in either of the following ways:

- **Identifying a command by its sequence number:** Each command in the history list is preceded by its sequence number in your command history. If you've displayed the history list and can see the command that you want to execute, you can re-execute a specific command by typing an exclamation mark followed by its sequence number. For example, using the history output displayed earlier in this section, the command `!236` would re-execute the `vi ~/.bashrc` command.

- **Identifying a command by name:** You can re-execute the most recently executed command that begins with a certain command name by typing an exclamation mark followed by the name of the command that you want to re-execute. For example, using the history output displayed earlier in this section, the command `!vi` would re-execute the `vi ~/.bashrc` command.

TIP **If you want to redisplay a previous command and edit it before re-executing it, you can use the history mechanism's `:p` modifier to simply print the command without executing it. You can then type `Ctrl+P` to redisplay that command, edit it however you wish, and then press Return to execute the modified command.**

If you've already executed several similar commands, you can narrow down the list of those shown by the history command by displaying only commands that contain a certain substring. For example, the command `history | grep vi` uses a pipe and the `grep` command to display only lines of output from the history command that contain the string `vi`. (See the section "Pipes and Input and Output Redirection" later in this chapter for more information about connecting different commands in this way.)

The bash shell's command-line history mechanism is extremely powerful, and this discussion only scratches the surface of the myriad ways that you can save yourself typing by recalling and updating previously executed commands. For all of the gory details on the shell's command-line history mechanism, use the `info bash` command to display the current online help for the bash command.

Using Command and Filename Completion

As you've seen throughout this chapter, the command line provides a very powerful, flexible mechanism for executing commands, but one of the goals of all UNIX and Linux developers seems to have been to allow users to type as few characters as possible in order to execute commands. This has given us cryptic commands like `ls` and `pwd`, but has probably saved wear and tear on thousands of keyboards over the years.

This same "type as little as possible" philosophy extends to the shell in a variety of ways. One of the most useful is the use of the tab key to help the shell match and display commands and filenames. Pressing the tab key while typing the name of a command will cause the shell to expand the name of that command as uniquely as possible. For example, because I use the Emacs text editor for almost all activities related to text files (see Chapter 15, "Working with Text Files on Ubuntu," for a discussion of my love affair with Emacs), I type commands such as `emacs this_is_the_name_of_a_file.txt` a few hundred times a day. In true Linux fashion, I can save myself a few characters of typing by typing `ema` on the command line, and then pressing the Tab key. This displays the first matching command, which happens to be `emacs`, and I can then press a space and enter the name of the file that I want to edit.

But wait, there's more! The Tab key will also match and insert the names of filenames. For example, if I were indeed trying to edit a file called `this_is_the_name_of_a_file.txt`, I could type the first few characters of its name and press the Tab key to (hopefully) match and insert the complete filename so that all I'd have to do is press return to begin editing the file. This is a more common use of the Tab key than completing the names of commands because, as we've seen throughout this chapter, the names of Linux commands tend to be short. On the other hand, they are cryptic, so if you don't remember the exact spelling of the command that you want to execute, the Tab key can often help you figure it out.

Pressing the Tab key matches as much of the name of a command or file as it can. If multiple commands or filenames match what you've typed before you press the Tab key, pressing Tab will have no effect because the shell can't tell which of the possible commands or filenames you mean. In this case, you can press the Tab key twice (quickly) to cause the shell to list all of the available matches. You can then type the additional few characters that are required to uniquely identify the command or filename that you're actually interested in.

Using Wildcards

Another example of the "type as little as possible" philosophy is the ability of the shell to identify groups of files whose names match certain patterns. The characters used to define those patterns are known as *wildcards* because they can match any single or multiple group of characters. The shell supports four primary types of wildcards:

- ?: (question mark) which matches any single character
- *: (asterisk) which matches any group of characters
- `{string1, string2}`: (curly braces) which match either string1 or string2
- `[char1,char2,...]`: (square brackets) which match any character or range of characters inside the square brackets

These wildcards can be extremely handy when working with directories containing multiple files with similar names. Consider the following directory containing a number of audio files in different formats:

```
$ ls
d1t01.flac  d1t04.flac  d1t07.flac  d1t10.flac  d2t02.flac
d1t01.mp3   d1t04.mp3   d1t07.mp3   d1t10.mp3   d2t02.mp3
d1t01.wav   d1t04.wav   d1t07.wav   d1t10.wav   d2t02.wav
d1t02.flac  d1t05.flac  d1t08.flac  d1t11.flac  d2t03.flac
d1t02.mp3   d1t05.mp3   d1t08.mp3   d1t11.mp3   d2t03.mp3
d1t02.wav   d1t05.wav   d1t08.wav   d1t11.wav   d2t03.wav
d1t03.flac  d1t06.flac  d1t09.flac  d2t01.flac  Horslips info.txt
d1t03.mp3   d1t06.mp3   d1t09.mp3   d2t01.mp3   playlist.m3u
d1t03.wav   d1t06.wav   d1t09.wav   d2t01.wav
```

You'll note that these filenames follow a certain convention, which in this case is `disc <number> track <number>.format`. To list just the MP3 files in this directory that are the first track on a given disc, you could type the following:

```
$ ls d?t01.mp3
d1t01.mp3  d2t01.mp3
```

This uses the question mark to represent the disc number, which is a single digit in my sample filenames.

Similarly, to list all of the MP3 files in this directory, you could use the asterisk to represent any matching filename with the `.mp3` file extension, as in the following example:

```
$ ls *.mp3
d1t01.mp3  d1t04.mp3  d1t07.mp3  d1t10.mp3  d2t02.mp3
d1t02.mp3  d1t05.mp3  d1t08.mp3  d1t11.mp3  d2t03.mp3
d1t03.mp3  d1t06.mp3  d1t09.mp3  d2t01.mp3
```

Assume that you want to list just the files with extensions that are either `flac` or `wav`. You could do this using the curly brace wildcard in the following command:

```
$ ls *.{flac,wav}
d1t01.flac  d1t04.flac  d1t07.flac  d1t10.flac  d2t02.flac
d1t01.wav   d1t04.wav   d1t07.wav   d1t10.wav   d2t02.wav
d1t02.flac  d1t05.flac  d1t08.flac  d1t11.flac  d2t03.flac
d1t02.wav   d1t05.wav   d1t08.wav   d1t11.wav   d2t03.wav
d1t03.flac  d1t06.flac  d1t09.flac  d2t01.flac
d1t03.wav   d1t06.wav   d1t09.wav   d2t01.wav
```

Finally, suppose that you want to list all of the seventh and eighth songs in MP3 format from the first disc in this directory. You could do this using the square bracket wildcard, as in the following example:

```
$ ls d1t0[78].mp3
d1t07.mp3  d1t08.mp3
```

You can combine any of these wildcards to identify increasingly specific sets of files. For example, the following would identify just the seventh and eighth songs from the first disk that are in `flac` and `wav` format:

```
$ ls d1t0[78].{flac,wav}
d1t07.flac  d1t07.wav  d1t08.flac  d1t08.wav
```

The bash shell provides a very powerful wildcard mechanism. Wildcards such as these are especially useful when you want to move, copy, or delete specific groups of files.

TIP **If you're familiar with wildcards on other types of systems, be careful: Linux wildcards may not work exactly the same way. For example, on some systems, you can type something like `ren *.txt *.foo` to rename groups of files with the `.txt` extension to files with the same base filename but with the `.foo` extension. This will not work on a Linux system because wildcards are expended before the command is executed, and the `*.foo` portion of the command line wouldn't match anything because no files with that extension currently exist.**

Pipes and Input and Output Redirection

As discussed in Chapter 1, a core element of the basic philosophy of UNIX, and therefore Linux, is that utilities should do one thing and do it well. This obviates the need for a lot of duplicate code by enabling one command to process the output produced by another command. For example, you might want to find the number of files in a directory by listing the files in a directory, and then counting the number of outlines. One easy way to do this is to write the output of the `ls` command to a temporary file, and then count the number of lines in that temporary file using the Linux `wc -l` command. Linux makes it easy to write the output of a command to a file by using what is called output redirection, which is represented on the command line by the greater-than symbol (>). As an example, the following command writes the output of the `ls` command to the temporary file `/tmp/my_dir_list.txt`:

```
$ ls > /tmp/my_dir_list.txt
```

I could then count the number of lines in that file using the `wc -l` command, as in the following example:

```
$ wc -l /tmp/my_dir_list.txt
8
```

There are apparently eight files in my current directory. However, using a temporary file required some extra typing and also requires that you clean up after yourself. Luckily, Linux comes to the rescue again with the notion of using a special symbol, the vertical bar, to connect the output of one command to the input of another. This symbol (|) is known as a *pipe*. You could use a pipe to avoid the temporary file and get the same results by typing a command such as the following:

```
$ ls | wc -l
8
```

If you find yourself performing many command-line operations, you will find both the output redirection and pipe concepts to be very convenient. Redirecting output is an easy way of keeping track of the contents of files and directories, recording system information, and so on.

An even more interesting use of output redirection is its ability to append some amount of information to an existing file. In the previous example, using a single greater-than symbol creates a file if it doesn't exist, but overwrites the contents of that file if it already exists. You may want to append to an existing file rather than completely overwrite it, which you can easily do by using two greater-than signs rather than a single one, as in the following example:

```
$ ls > /tmp/my_dir_list.txt
$ ls >> /tmp/my_dir_list.txt
$ wc -l /tmp/my_dir_list.txt
16
```

As you can see, the temporary file now contains twice as many lines, which is what you'd expect because you wrote the same eight lines to it twice.

The bash shell also supports input redirection, which is the ability to read data from a file and use that file as input to another command. For complete information on input redirection and redirection in general, see the online help for the bash shell. You'll find that you use output redirection and pipes much more often than input redirection, but you never know …

Introducing Job Control

Throughout this chapter, you've always executed commands at the shell prompt, and then waited for them to finish. That's fine for commands that are relatively quick, but what about commands that run for a long time? The bash shell provides a nice solution for this sort of thing by enabling you to start commands but run them in what is known as the background. This enables you to type a command, indicate that it should run in the background, and that command will continue to execute even though the shell prompt is redisplayed in your current shell window. For example, suppose that you want to use the Emacs text editor to edit the file `myfile.txt`, but also want to be able to use your current `xterm` or GNOME Terminal window while Emacs was running in its own window on the screen. To start a command in the background, you simply follow its name with an ampersand, as in the following example:

```
$ emacs myfile.txt &
[1] 7539
```

Starting a command in the background is known in Linux-speak as starting a background job, and causes the shell to display two bits of information:

- The job number (inside square brackets)
- The process ID of the process that was started to run your command

The job number is useful because it is possible to run multiple commands in the background at the same time. The job number therefore provides a unique way of identifying any particular background command. The job number is an identifier for that background job within the context of your current shell. The process identifier is a unique identifier for that process in terms of your entire Ubuntu system.

Why care about identifying these processes? Primarily so that you can either terminate them or wake them up if something goes wrong. For example, suppose that you accidentally started a command in the background, but that command actually requires interactive input. To bring a background command back to the foreground (meaning that it is interactively running in the current shell), you can simply refer to it by using a percent symbol and its job number, as in the following example:

```
$ %1
```

This would bring the `emacs` command back to the foreground so that anything you typed in the shell is actively sent to the command. Suppose, however, that you've lost your Emacs window on the screen and simply want to terminate the background job manually. You can do this by using the Linux `kill` command, providing the job number to identify the job that you want to terminate, as in the following example:

```
$ kill %1
```

Knowing not only the job number for a background process but its system-wide process identifier enables you to use the `kill` command to terminate a process from another shell. Because job numbers are unique to a specific shell, you can use the `kill` command to terminate a process-by-process ID from any shell or terminal session, as in the following example:

```
$ kill 7539
```

Being able to execute long-running commands in the background is very important in today's multi-windowed graphical environments because you usually want to be able to continue to use your current shell while the long-running command continues on its merry way.

As always, job control is much more flexible and powerful than the introduction given here describes. For complete information on job control, use the `info` command to see the online reference information for the bash shell.

Exploring the Bash Configuration File

When you start a bash shell on your Ubuntu system, it reads a certain amount of configuration information from various text files on your system. In the order they're read, these configuration files are the following:

- `/etc/profile` (login shells only)
- `/etc/bash.bashrc` (all bash shells)
- `~/.bash_profile` (login shells only)
- `~/.bashrc` (all bash shells)

NOTE **The ~ symbol is a bash shortcut for referring to any user's home directory. `~/.bashrc` is the file `.bashrc` in my home directory; `~juser/.bashrc` is a reference to the user called juser's home directory.**

The first two of these are shell initialization files that are read and incorporated when any user logs in on the system (`/etc/profile`) or starts a bash shell (`/etc/bash.bashrc`). The second two are personal configuration files that are stored in each user's home directory. As with the system-wide configuration files, the `~/.bash_profile` file is read by the user's login shell, while the `~/.bashrc` file is read by any shell that the user creates, such as in an `xterm` or GNOME Terminal window. Of these, the `.bashrc` file is most important for two reasons:

- The default `~/.bash_profile` on Ubuntu systems reads it even in the context of a login shell.
- Starting a bash shell is a much more common operation than logging in. A user only logs in once, though he or she may subsequently create many different `xterm`s or GNOME Terminal windows.

The `.bashrc` file enables each user to customize various aspects of the bash shell, such as where it looks for commands, the prompt that it displays, the number of commands that it remembers in your command-line history, and many more things. The `.bashrc` file enables you to create your own shortcuts to

specific Linux commands if you still feel that you're having to type too many characters and even enables you to create your custom shell commands. Lines in any bash configuration file (or shell script, for that matter) that begin with the hash mark (#) are considered to be comments and are not interpreted.

TIP **All Linux shells provide a command language and constructs for looping, and support conditional expressions to make it easy to write shell scripts. Explaining shell programming is outside the scope of this introduction, but there are many excellent books and Web sites on that topic. Two of my favorite Web sites for this sort of information are the Linux Shell Scripting Tutorial at** `www.freeos.com/guides/lsst/` **and the Linux Documentation project's Advanced Shell Scripting Guide at** `www.tldp.org/LDP/abs/html`.

When a user account is created on your Ubuntu system, the default `.bashrc` and `.bash_profile` files are created by copying them from the templates with the same names that are stored in the directory `/etc/skel`. The following is the beginning of the default `.bashrc` file used for all Ubuntu user accounts:

```
# ~/.bashrc: executed by bash(1) for non-login shells.
#
# If not running interactively, don't do anything
[ -z "$PS1" ] && return
[much more text deleted]
```

The last statement in this example stops shells from reading the default `.bashrc` file if they are not running in interactive mode. It does this by using the `[` command to test the value of the PS1 environment variable, which defines the user's prompt, and to return to the parent shell if this variable is not set (which it would not be in the context of a shell used to run a shell script).

What's an environment variable? I'm glad you asked.

Using Environment Variables

Environment variables are variables that are defined in the context of a shell and which are read and used by applications on your system, including the shell itself. One of the most commonly used environment variables is the PATH environment variable, which contains a colon-separated list of directories that the shell searches (in order) to try to find commands. You can examine the environment variables that are defined in a shell using any of three mechanisms:

- Using the shell's built-in `set` command to display all shell variables that are currently defined
- Using the `printenv` command to display a list of selected (or all) environment variables that are currently defined
- Using the shell's built-in `echo` command to display the value of a specific environment variable

Both the `set` and `printenv` command produce large amounts of output, so the `echo` command is most commonly used to display the value of a single environment variable, as in the following example:

```
$ echo $PATH
/home/wvh/bin:/usr/local/Adobe/Acrobat7.0/bin:\
/usr/local/gcc_svn/bin:/home/wvh/BitTorrent/src/BitTorrent:\
/home/wvh/cxoffice/bin:/usr/lib/mit/bin:/usr/lib/mit/sbin:\
/usr/NX/bin:/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:\
/sbin:/bin:/usr/bin/X11:/usr/games:
```

NOTE **The value of the PATH environment variable is a single string containing colon-separated directories. This example has been split across multiple lines for formatting purposes.**

When using the `echo` command to display the value of an environment variable, you must refer to the environment variable by preceding its name with a dollar sign (`$`). This tells the shell to echo the value of a variable rather than simply echoing a string, which is what would happen if you simply typed the following:

```
$ echo PATH
PATH
```

This example of the PATH environment variable has been heavily customized. The default setting for the PATH environment variable on Ubuntu systems is actually the following:

```
/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin:\
/usr/bin/X11:/usr/games
```

To add a directory to the list of directories that are searched for commands, you can execute something like the following from a shell prompt:

```
$ export PATH=/usr/wvh/bin:$PATH
```

This command sets the environment variable and exports it to the general shell environment. This example adds my personal bin directory as the first element of my PATH, so that the shell looks there first for any commands that I try to execute. To make this change permanent, I would add this line to the end of my `~/.bashrc` file.

The `echo` command is useful for displaying the value of a single environment variable. You can also use the `printenv` command to do the same thing, as in the following example:

```
$ printenv PATH
/home/wvh/bin:/usr/local/Adobe/Acrobat7.0/bin:\
/usr/local/gcc_svn/bin:/home/wvh/BitTorrent/src/BitTorrent:\
/home/wvh/cxoffice/bin:/usr/lib/mit/bin:/usr/lib/mit/sbin:\
/usr/NX/bin:/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:\
/sbin:/bin:/usr/bin/X11:/usr/games:
```

Because the `printenv` command is designed to display environment variables, you do not need to precede the name of the environment variable with a dollar sign.

The `printenv` and `set` commands are more commonly used to display the values of all environment variables that are currently defined. You can then search for a certain substring by piping their output to the `grep` command, as in the following example:

```
$ printenv | grep HIST
HISTCONTROL=ignoreboth
HISTFILE=/home/wvh/.bash_history
HISTFILESIZE=1000
HISTSIZE=1000
```

This sample command shows me the name and value of all environment variables that contain the string HIST, which are environment variables related to the shell's command-line history mechanism (discussed earlier in this chapter).

Defining and Using Aliases

Another interesting feature of the bash shell is its ability to enable you to define aliases for various commands. An alias is a string that the shell recognizes and expands into some other command. Some of the most common aliases used on Linux systems are those for the `cp` and `rm` commands that automatically execute those commands with their interactive option, which prevents you from accidentally overwriting or deleting files that you didn't want to overwrite or delete by prompting you to confirm each operation. These aliases are defined in the following way at the bash command line (or, permanently, in your `~/.bashrc` file):

```
$ alias cp="cp -i"
$ alias rm="rm -i"
```

After defining these aliases, the shell automatically replaces either the `cp` or `rm` commands with the command that it has been aliased to.

To simplify defining aliases (and to minimize the number of changes that you make to your `~/.bashrc` file), the default `~/.bashrc` file used on Ubuntu system contains the following entries, which are initially commented out:

```
#if [ -f ~/.bash_aliases ]; then
#    . ~/.bash_aliases
#fi
```

If uncommented by removing the leading hash marks, this clause enables you to define all of your aliases in a file called `~/.bash_aliases`, which would be located in your home directory.

Summary

This chapter introduced you to the Linux command line, the terminology used when discussing Linux command-line interpreters, and the basic capabilities of the bash shell, the command-line interpreter that is used by default on Ubuntu Linux systems. It discussed how to access a shell on your Ubuntu system and some of the more common commands that most people execute within a command-line environment. It concluded with an explanation of the configuration files used by the bash shell and some of the more common items that they contain.

Hand in hand with the idea of working from a command line is the notion of working with text files, which is discussed in Chapter 15. Though many word processors store their files in special binary formats nowadays, text files are alive and well on the Internet and on Linux systems. Documents written using markup languages such as HTML, SGML, and XML are all text files that are easily created and edited in a standard text editor. Most Linux servers, including those discussed in the last few chapters of this book, are configured by setting or changing operational parameters that are stored in text-format configuration files. For these reasons and many more, learning how to use a text editor is a fundamental Linux skill.

Chapter 15

Working with Text Files on Ubuntu

Most of us are used to working with files in various application-specific formats, often identified by their file extension or a special icon on our graphical desktops. We're all familiar with the doc files produced by graphical word processors such as Microsoft Word, pdf files used by Adobe Acrobat and other PDF readers, xls files produced by spreadsheets such as Excel, fm files produced by FrameMaker, ppt files produced by PowerPoint, and so on. All of these files contain application data in a specific, often binary, format that lets the associated application make the best possible use of these files, but which often makes them hard to use with any application other than the one that created them.

At the other end of the spectrum from these application-specific files is the lowest-common-denominator file format known as text files or, to be a bit more specific, *ASCII text* files. The contents of these files consist of only the standard letters, numbers, punctuation, and symbols that you find on a computer keyboard. The nice thing about these types of files is that they are easy to read, easy to process, and easy to work with in general.

Regardless of what you use your Ubuntu Linux system for, you will almost certainly end up editing a text file one of these days. If you write code, it is almost certainly in standard ASCII text files because they provide a lowest-common-denominator format that can be processed by any compiler. If you're a writer, you may want to use text files because you can work with them on any system, with any editor or word processor — so you can work on them just as easily on your Ubuntu Linux system as you can on your Mac OS X system.

An even more critical point is that most system-level Linux utilities use one or more files to hold their configuration and initialization information. Even though most Linux utilities can also be configured using a graphical interface, editing a text file is often faster and simpler than locating and starting the right graphical tool, finding the correct tab or panel, and so on.

This chapter begins by briefly describing the history of editors on UNIX systems and the religious war over the "right" approach to editing text that still

IN THIS CHAPTER

- Using popular editors for UNIX-like systems
- Using vi
- Using Emacs
- Using the graphical gedit editor
- Using the graphical KEdit editor
- Using the graphical KWrite editor
- Using the graphical Kate editor
- Using other editors for Linux systems

rages on in Linux/UNIX circles today. It then provides detailed information about using vi and Emacs, the two primary editors used on Linux systems, and includes an introduction to ed, a fast, simple, and minimal editor that is perfect for quick text file changes. The chapter concludes by discussing gedit, a graphical text editor provided in the GNOME environment that walks the line between pure text and graphical editing, and a variety of similar editors for the KDE desktop environment.

Introduction to Linux Text Editors

The first versions of UNIX for the PPD-11 that were distributed by Bell Telephone Laboratories included a small, line-oriented text editor called ed, which was written by Ken Thompson. Various people enhanced ed, most notably George Coulouris at Queen Mary College (QMC) in the United Kingdom, who added interactive support for editing single lines as part of his editor for mortals (em) editor, which he brought with him to the University of California at Berkeley (UCB) in 1976. At the same time, Ken Thompson had gone to UCB to teach for a year or so, and while there had begun working with Bill Joy and Chuck Haley on the UNIX enhancements that would eventually make up the original Berkeley Standard Distribution (BSD).

Haley and Joy hacked em for a while, revising it into en and then a bunch of other names, working their way through the alphabet until the arrived at ex, which was still line-oriented. In 1976, however, Joy created vi (which stands for visual), a screen-oriented editor that shared the ex command set, thanks to motivation provided by a number of ADM-3A terminals which were donated to UCB. The flexibility provided by vi in terms of using the entire screen, thanks to the `/etc/termcap` (terminal capabilities) file, was a huge step in divorcing detailed knowledge about the hardware from an editor. Editing under UNIX has never been the same.

Around the same time, on the other side of the country, a bunch of Lisp and ITS/TOPS-20/TWENEX hackers at the MIT Artificial Intelligence (AI) Lab were enhancing an aging line editor for their PDP-10 systems known as Text Editor and Corrector (TECO). With some inspiration from graphical editing work being done at Stanford's AI Lab, Richard Stallman (later of GNU and FSF fame) and Guy Steele collected and developed a version of TECO to which users could add their own commands. They also created and aggregated a set of TECO commands into a single package that provided support for screen-oriented editing using TECO. The TECO commands consisted of a set of macros (sequences of commands that execute together), so the set of macros became known as Editing MACroS, or more simply, EMACS.

This was all well and good if you happened to be using a DEC PDP-10, and Emacs would be an interesting footnote were it not for the efforts of Dave Conroy, who wrote a version of Emacs called MicroEmacs that ran on just about anything, and James Gosling (later of Java fame), who wrote a version of Emacs for UNIX systems in 1981. Unlike the Emacs macros that were their inspiration, these were actual compiled programs that ran directly on specific systems as compiled binaries. Gosling's version of Emacs, cleverly named Gosling emacs or simply gosmacs, served as the original foundation for an open source version of Emacs for UNIX. Richard Stallman began writing GNU emacs in 1984. For legal reasons, Stallman eventually replaced all of the Gosling emacs code in GNU emacs, and replaced Gosling's MockLisp emacs configuration language with a full-blown internal Lisp interpreter to parse configuration files and other startup data.

NOTE **Over the years, several alternative expansions of the EMACS acronym have been offered to replace the original Editing MACroS. Some of my favorites are "Escape Meta Alt Control Shift," referring to the heavy dependence of Emacs on the Ctrl and Escape keys in its command set, and "Eight Megabytes and Constantly Swapping," which was funnier when 8MB was a lot of memory for a computer, and referred to its large memory requirements (for the time). I think the garage door opener I use today has 8MB of memory.**

Modal vs. Modeless Editors

As mentioned earlier in this chapter, among Linux and UNIX users, there is no more entertaining or more long-running flame war than the debate over which editor is better: vi or Emacs. If all of the characters expended on this topic alone had been given instead to random simian typists, they would have indeed produced the works of Shakespeare. As a new Ubuntu fan, you'll eventually hear something about this, so this sidebar gives you a little background.

The basic difference between how you use vi and Emacs boils down to the following: Emacs is what is known as a *modeless editor*, which means that any standard characters that you type are always inserted into whatever you're editing (which is therefore known as *insert mode*), whereas vi is what is known as a *modal editor* in which insert mode is just one available mode. Both Emacs and vi are extremely configurable, although vi doesn't provide an internal implementation of Lisp to "simplify" creating your own functions. On the other hand, because vi doesn't provide an internal implementation of Lisp, it is much smaller and faster to execute than Emacs. Emacs is more powerful because it includes a mail client, newsgroup reader, and even a few games. On the other hand, you may just be looking for a text editor, not a one-stop-shopping computing environment.

People often complain that vi and Emacs are difficult for new users to learn. In my experience, it just depends where you want new users to get confused. In Emacs, the chances of figuring out that you hold down the Ctrl key, and then press a particular two-key combination to save your file and exit are pretty small. However, vi rarely does anything useful until you enter insert mode, and trying to get out of insert mode and figuring out how to save your file has brought many a Linux/UNIX novice to tears.

Here's my take: I have used versions of Emacs since gosmacs, and I've written this book in Emacs. My fingers work that way. However, I use vi regularly when I'm performing most system tasks and especially when I'm sitting down at someone else's machine for the first time. Why? Because vi starts more quickly than Emacs, and because Emacs fans often extensively customize their Emacs key bindings. The key sequence that justifies a paragraph on my system might post a random snapshot to the alt.binaries newsgroup if I type it on your system. Your mileage may vary.

If you're interested in more details on the history of editing text on UNIX systems (and thus on Linux), here are some great links:

- ed History (and more): `www.english.uga.edu/hc/unixhistory.html` (from Peter H. Salus, *A Quarter Century of UNIX*; Reading, MA: Addison-Wesley, 1994).
- George Colouris on em: `www.dcs.qmul.ac.uk/~george/history/`.
- Bill Joy Interview: `www.cs.pdx.edu/~kirkenda/joy84.html` (originally printed in the August 1984 issue of *Unix Review* magazine).
- The Lemacs/FSFmacs Schism: `www.jwz.org/doc/lemacs.html`.

Using vi

Because the actual source code for the "real" version of vi is owned by AT&T, vi fans have completely rewritten vi from the ground up. Three popular vi clones are available today: elvis, nvi (new vi), and vim (vi improved). The version of vi used on Linux systems is vim, but nvi is quite popular and is the default

version of vi used on many open source BSD-like systems such as OpenBSD, FreeBSD, and so on. Elvis is also a fine clone of vi, supporting multi-file and multi-window editing, as well as X11 support. However, the discussion of vi in this book focuses on vim because that's what you get by default with Ubuntu and Kubuntu installation. Therefore, I'll generally refer to vim as vi throughout this section because it is indeed compatible with vi, and the package that installs vim on your Ubuntu system also installs an alias for vim called vi — and vim is, after all, just vi, *improved*. In any sections that are specific to vim and do not apply to vi in general, I'll refer to vim explicitly.

TIP **The default version of vim that is installed on Ubuntu and Kubuntu systems is a small package called `vim-tiny`. After all, not everyone is a vi/vim fan! However, if you want to explore the full power of vim and work on your basic UNIX/Linux chops, you should install the `vim` package, which will also install another package called `vim-runtime`. Together, these packages provide many enhancements to vim. For example, the `vim-runtime` package includes an application called `vim-tutor`, which starts vim with a copy of a tutorial on using vim. If you want to actually walk through many examples of using vi and get a feel for its command set, `vimtutor` is an excellent starting point.**

As discussed in the previous section, vi is a modal editor (which, if you skipped the previous section, means that the same keys do different things depending on the mode that you're in). The primary vi modes are as follows:

- **Normal:** vi is in normal mode when you first start vi. In this mode you can enter all of the standard cursor movement and text deletion commands, but not enter text into the file that you are editing. This is the primary source of confusion for new vi users because most people expect to start an editor and begin editing text.
- **Insert:** In this mode, typing the standard alphanumeric and symbol keys on your keyboard enters those characters into the file that you are editing. Insert mode is the mode that most people expect to be in when you start a text editor. The fact that once in insert mode you need to press a special key (the Escape key) to return to command mode is the second primary source of confusion for new vi users who find themselves trapped in insert mode.
- **Command-line:** In this mode, the cursor is positioned at the bottom line of your vi window, and you can enter a single line that is interpreted as a vi command, search request, or a request to invoke an external command on a portion of the file that you are editing (often referred to as a *filter* command).

The modal editor concept may initially seem confusing, but it makes vi a fast, extremely powerful editor. Once you learn how to get from one mode to another and work with vi for a while, you will find it to be an excellent editor. The fact that some version of vi is found on every UNIX-like system makes it well worth your time to learn.

One of the coolest features of vi is the ability to preface most vi commands with a numeric argument, movement, or scoping command, which executes that single command the specified number of times, the number of times specified by the movement command that you've specified, or applies a command to a specified portion of your file. I'll provide examples of this throughout the rest of this section, as well as some of my favorite vi tips and tricks that I've accumulated over the years.

NOTE **All of the examples in the vi portions of this chapter show vi in an `xterm` window rather than the GNOME Terminal to emphasize that vi is a command-line application that does not provide or use menus by default. For information on graphical versions of vi that do provide and use menus, see "Graphical Versions of vi" later in this chapter.**

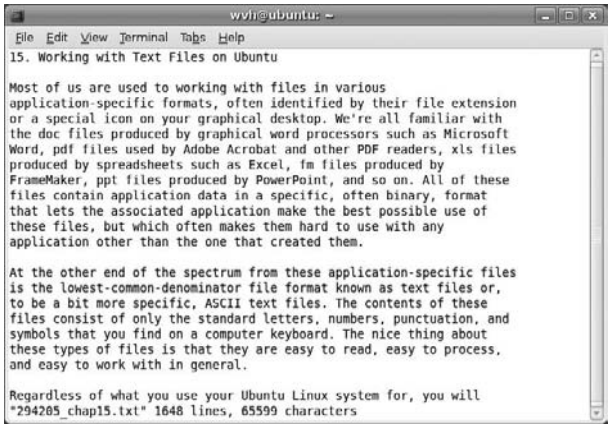
Starting and Exiting vi

The core tasks for any application are how to start it and how to get out of it. As mentioned previously, getting out of vi the first time can be tricky because of its modal nature, but after that, it becomes second nature.

Because vi is a command-line application, you usually start it by simply typing the vi command in a GNOME Terminal or xterm window. You can optionally provide the name of the file that you want to edit on the command line, and vi will automatically open and display that file for you. Figure 15-1 shows vi editing the text of this chapter in an xterm window.

FIGURE 15-1

An xterm window showing the vi editor



You can exit vi and return to the command prompt in several ways, but you first have to be in vi's normal mode to do so. If you are editing a file and want to return to normal mode, simply press Escape. If the key sequence ^[displays in the file that you're editing, press Escape again. You can then exit from vi in the following ways:

- To exit unconditionally without saving any changes, type the :q! or ZQ command. The :q command is the standard quit command, which will complain if you have made changes to the current file and have not saved them. The :q! command essentially means "Yes, I really want to quit and I know what I'm doing."
- To save any changes that you've made to the current file and exit, type the :x or :wq command. The :x command is specific to vim, while the :wq command is the traditional vi exit mechanism (it stands for "write and then quit"). You can also write any changes to the current file and exit by typing the ZZ command (no colon required).
- To specify the name of the file that you want to save your changes to, type the :w *filename* command, where *filename* is the name of the file that you want to create. If the file already exists, you will have to use the :w! *filename* command which, loosely translated, means "Yes, I really want to write this file and I know what I'm doing."

Inserting Text in vi

Once you've started vi, you'll certainly want to start creating the program, letter, or other document that you're working on. As mentioned in the previous section, starting vi leaves you in normal mode, where you can move the cursor and enter commands, but not actually enter text.

To begin inserting text, type the i command. This causes vi to enter insert mode. Any characters that you type will be inserted into the file that you are editing until you exit from insert mode. As a reminder, the easiest way to exit from insert mode is by pressing the Escape key on your keyboard.

NOTE **All commands in vi and its clones are case-sensitive.**

Using the i command begins inserting text at the current cursor position, which can be awkward. If you want to begin adding text immediately after the current cursor position, use the a command. To begin appending text after the end of the current line, use the A command. The a and A are just specialized commands for entering insert mode at a specific point — to exit insert mode after using these commands, press the Escape key.

Moving Around in vi

The vi editor is knee-deep in cursor movement commands. The most common are as follows:

- h or left-arrow: Move the cursor backward one character in the current line.
- l or right-arrow: Move the cursor forward one character in the current line.
- j or down-arrow: Move the cursor down one line, maintaining the current column position in the new line when possible.
- k or up-arrow: Move the cursor up one line, maintaining the current column position in the new line when possible.

Cursor movement commands such as h, j, k, and l may initially seem pretty strange or at least counterintuitive, but remember that vi was originally designed to be used on terminals, some of which didn't have the arrow keys. Ask your grandfather.

NOTE **When in insert mode, the old-time vi cursor movement commands (`h`, `j`, `k`, and `l`) insert those letters rather than moving the cursor. (vi is a modal editor, remember?) However, the arrow cursor movement keys "do the right thing" when you're in insert mode — they simply move the cursor but leave you in insert mode.**

As mentioned previously, you can repeat these simple cursor movement commands a specific number of times by prefacing them with a numeric argument. For example, you can move the cursor forward five letters by typing 5l, up five lines by typing 5k, and so on. This is handy, but requires that you count how far you want to move the cursor first. That's somewhat tedious when you really want to do things like move to the next word, move to the previous word, or (more commonly) go to some logical place in the current text. To simplify things, vi provides handy cursor movement commands like the following:

- b: Move backward one word.
- w: Move forward one word.
- ^: Move to the beginning of the current line.
- $: Move to the end of the current line.

- (: Move to the beginning of the current sentence.
-): Move to the end of the current sentence.
- {: Move to the beginning of the current paragraph.
- }: Move to the end of the current paragraph.
- 1G: Go to the first line in the current file.
- 0G: Go to the last line of the current file. Just pressing G in vim does the same thing, but that may not work in all versions of vi.

The last two commands are actually examples of the standard vi commands to go to a specific line, and they highlight how nicely vi commands work with numeric arguments. For example, you could move to the 14th line in your file by typing the 14G command. Similarly, you could move to the end of the third paragraph from your current position by typing the 3} command.

In addition to these cursor movement commands, vi also provides several commands to enable you to scroll the screen without moving the current cursor position (unless you want to, of course). These screen movement commands are as follows:

- Ctrl+D: (Hold down the Ctrl key and press the letter d.) Scroll down one screen in the current file.
- Ctrl+E: (Hold down the Ctrl key and press the letter e.) Scroll down one line in the current file.
- Ctrl+U: (Hold down the Ctrl key and press the letter u.) Scroll up one screen in the current file.
- Ctrl+Y: (Hold down the Ctrl key and press the letter y.) Scroll up one line in the current file.

The Ctrl+D and Ctrl+U commands scroll the screen but leave the cursor at the same position on the screen (though the text underneath the cursor has changed). The Ctrl+E and Ctrl+Y commands keep the cursor at the same position in your text until that line scrolls off the screen, at which point the current position is preserved but the text below it has changed.

Deleting and Changing Text in vi

As most operations, you can delete text in vi with several different commands. To delete the character at the current cursor position, press the letter x. You can, of course, delete a specific number of characters by prefacing the x command with a numeric argument. For example, 14x deletes 14 characters from the current character position.

Although the x command is handy when deleting a few characters, most people think of deletion in terms like “delete the current word,” “delete everything to the end of the current line,” or “delete the rest of the current paragraph.” To perform these sorts of actions, you use the d command, followed by a cursor movement command that delimits the end of the area that should be affected by the delete operation. For example, the command dw deletes all of the text from the current cursor position to the end of the current word, d$ deletes everything from the current cursor position to the end of the current line, and d} deletes everything from current cursor position to the end of the current paragraph. If you remember, the second character of each of these commands is one of the cursor movement commands described in the previous section.

When you’re editing a file, it is common to delete some text and then immediately replace it with new text. You could do that with a combination of one of the delete commands, followed by an insert or

append command. To save typing, vi provides the c (change) command to do both of these things (deleting a delimited portion of your text and entering insert mode) at the same time. For example, you could use the cw command to change the text from the current cursor position to the end of the current word, the c$ command to change the text from the current cursor position to the end of the current line, and so on. When using vi's change command, you are placed in insert mode, so you must use the Escape key to return to normal mode when you're done entering text.

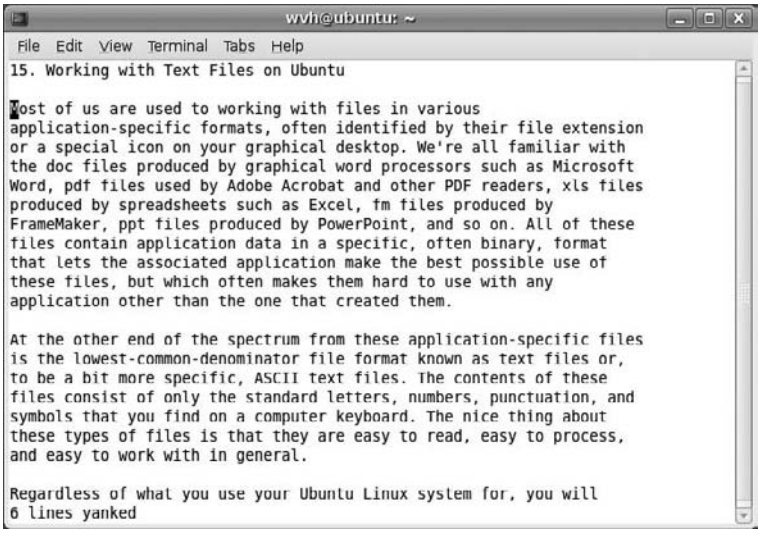
Cutting, Copying, and Pasting in vi

It's easy enough to cut and paste text using a graphical text editor because a mouse-aware text editor enables you to copy and cut text with a combination of the mouse and special keyboard commands. You can, of course, do the same thing in vi, but it requires a bit more knowledge of vi internals and visualization.

Cutting text is done using the delete command (explained in the previous section), which actually just puts the text in an internal storage area (known as a buffer) from which you can subsequently paste it back in. However, the vi editor also provides the y (yank) command to copy a specified region to a buffer so that you can subsequently paste it somewhere else. As with the delete command, you can delimit the area that you want to copy by following the y command with a cursor movement command. For example, y$ copies the text from the current cursor position to the end of the line into the buffer, while y} copies all of the text from the current cursor position to the end of the paragraph into the buffer. If you copy or delete anything more than a single line, the number of lines yanked into the buffer is displayed at the bottom of your vi window. Figure 15-2 shows the message displayed by vi running in an xterm window after copying the current paragraph (which is six lines long) into the buffer.

FIGURE 15-2

The vi editor displaying a yank status message



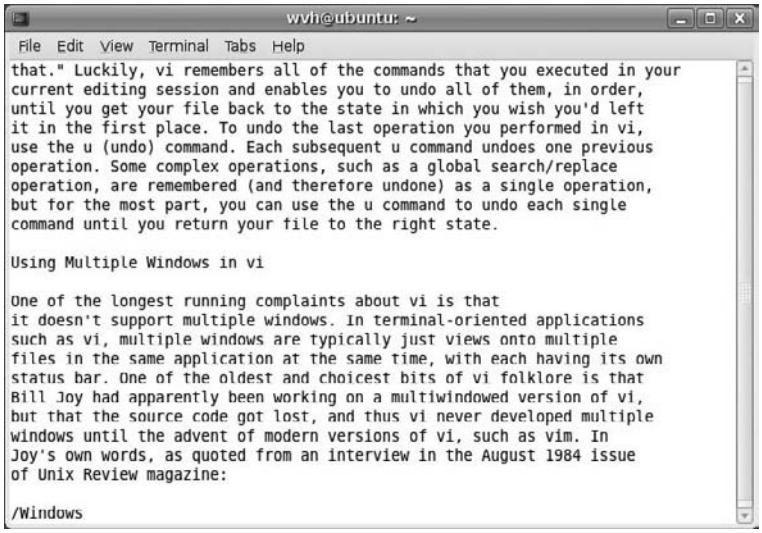
Of course, cutting or copying text isn't all that interesting unless you can paste it back in somewhere, for which vi provides the p and P commands. To paste a single line from the buffer back into your text at the current cursor position, use the p command. However, if you have copied multiple lines into the buffer, you should preface the p command with the number of lines that you want to paste back in. If you didn't happen to notice how many lines you yanked and you simply want to paste all of the text that you recently copied to the buffer, you can simply use the P command to paste it all back in at once.

Searching for and Replacing Text in vi

The vi editor provides separate commands for searching forward and backward in your file. Pressing the / command, followed by the sequence of characters that you want to search for, and pressing return causes vi to search forward in the current file until (if) it locates a matching sequence of characters. Figure 15-3 shows vi running in an xterm window, getting ready to search for the string Windows in the text. If the string is found, vi adjusts the text display window and moves the cursor forward in the file to display the line of text containing the match. If no matching string is found in the text, vi displays a status message to that effect and does not move the cursor.

FIGURE 15-3

Searching for a string in vi



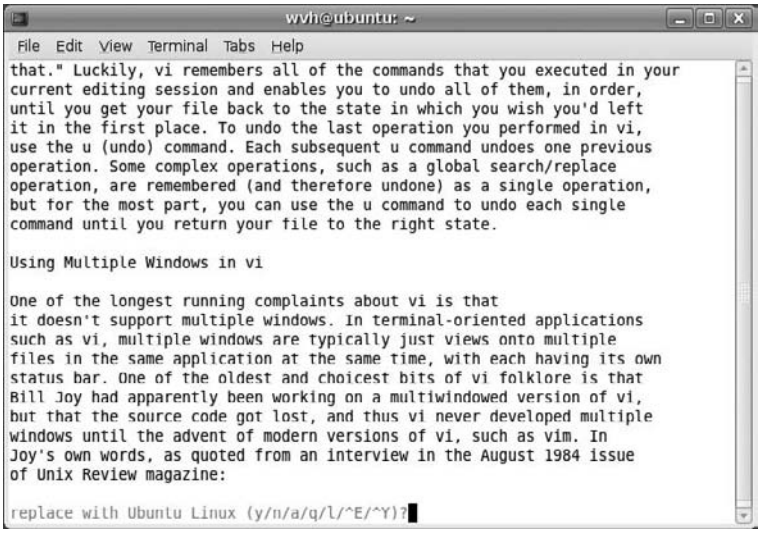
The ? command is the opposite of the / command, searching backward through the current file for a specified string. Both the / and ? commands are examples of vi's command-line mode, where pressing a command positions the cursor at the bottom of the screen, waiting for subsequent information.

The vi editor's command-line mode also provides a convenient text substitution command to enable you to search for a string and either globally or conditionally replace it with another. This is the :%s/foo/bar/ command, which (in this case) would search for the next instance of *foo* and replace it with *bar*. You can follow the final / with other characters to indicate whether the replacement should automatically be done

globally throughout the rest of the file (:%s/foo/bar/g), or whether vi should ask you for confirmation before performing each substitution (:%s/foo/bar/gc). Figure 15-4 shows a conditional replacement of the string Linux with Ubuntu Linux. Conditional replacements of this sort are often known as query/replace operations because they ask you for confirmation before doing each replacement.

FIGURE 15-4

A query/replace operation in vi



Note that the prompt at the bottom of the screen asks you to confirm (Y) or reject (N) each change, but also provides two cursor movement commands. The query/replace operation also supports the use of the Ctrl+E (scroll down one line) and Ctrl+Y (scroll up one line) commands to make it easier for you to see the context in which you're replacing the specified text.

Undoing Changes in vi

Unless you're a much better writer and programmer than I am, and therefore never make a mistake, you'll eventually find yourself making several changes to a file and then thinking, "Dang, I didn't want to do that." Luckily, vi remembers all of the commands that you executed in your current editing session and enables you to undo all of them, in order, until you get your file back to the state in which you wish you'd left it in the first place. To undo the last operation you performed in vi, use the u (undo) command. Each subsequent u command undoes one previous operation. Some complex operations, such as a global search/replace operation, are remembered (and therefore undone) as a single operation, but for the most part, you can use the u command to undo each single command until you return your file to the preferred state.

Using Multiple Windows in vi

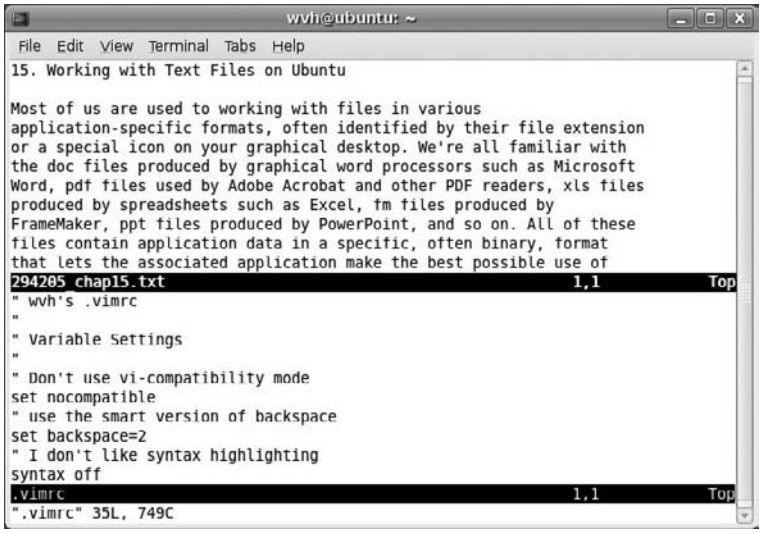
One of the longest running complaints about vi is that it doesn't support multiple windows. In terminal-oriented applications such as vi, multiple windows are typically just views onto multiple files in the same application at the same time, with each having its own status bar. One of the oldest and choicest bits of vi folklore is that Bill Joy had apparently been working on a multi-windowed version of vi, but that the source code got lost, and thus vi never developed multiple windows until the advent of modern versions of vi, such as vim. In Joy's own words, as quoted from an interview in the August 1984 issue of *Unix Review* magazine:

> …What actually happened was that I was in the process of adding multi-windows to vi when we installed our VAX, which would have been in December of '78. We didn't have any back-ups and the tape drive broke. I continued to work even without being able to do backups. And then the source code got scrunched and I didn't have a complete listing. I had almost rewritten all of the display code for windows, and that was when I gave up. After that, I went back to the previous version and just documented the code, finished the manual and closed it off. If that scrunch had not happened, vi would have multiple windows, and I might have put in some programmability — but I don't know…

Luckily, the vim folks have solved this problem by implementing several commands that enable you to split the current window, open another file in another part of your vim window, cycle between multiple files displayed in different windows, and so on. Figure 15-5 shows vim displaying two different files in separate windows within the same vim session.

FIGURE 15-5

Multiple windows in vim



The primary command-line commands for manipulating multiple windows in vim are as follows:

- :hide: Close current window (works only when multiple windows are displayed).
- :only: Close all windows except for the current one.
- :split filename: Split window horizontally and load filename in the new window.
- :sview filename: Same as split, but displays filename in read-only mode.
- :vsplit filename: Split window vertically and load filename in the new window.

TIP **You can start vim in multi-window mode by using the `-o` option and specifying multiple filenames on the command line, as in the following example:**

```
$ vim -o file1.txt file2.txt file3.txt
```

Once you've created multiple windows, you can move between and manipulate them using the following commands:

- Ctrl+W+up arrow: Move cursor up a window.
- Ctrl+W, Ctrl+W: Move cursor to another window (cycle).
- Ctrl+W_: Maximize current window.
- Ctrl+W=: Make all equal size.
- Ctrl+W+: Increase window size by one line.

After a bit of experimentation (and training your fingers to use a few new commands), working with multiple windows in vim is just as easy as working with them in any other text editor.

Customizing vim

By default, vim reads three configuration files when you first start the editor. These are as follows:

- The system configuration file /etc/vimrc.
- A per-user configuration file, located in each user's home directory and called .vimrc.
- The traditional per-user configuration file used by vi, located in each user's home directory and called .exrc.

It is not an error if any of these files do not exist, but together they provide a very flexible mechanism for customizing vim. Each of these files can contain any number of the following types of configuration commands:

- autocmd: Automatically execute specific commands based on the type of file that you are editing.
- iab: Define abbreviations that will be expanded into a specific sequence of characters when they are encountered as a word in the text that you are typing. This can function like the AutoCorrect feature in Microsoft Word.
- map: Map existing vim commands to a custom key or define your own vi commands.
- set: Set a vim variable to a value.

In addition to these specific commands, vim configuration files also support a rich set of conditionals that enable you to write your own functions and execute them in different contexts.

As an example of configuring or personalizing vim, my .vimrc file is as follows:

```
"
" wvh's .vimrc
"
" Variable Settings
"
" Don't use vi-compatibility mode
set nocompatible
" use the smart version of backspace
set backspace=2
" I don't like syntax highlighting
syntax off
" define the point at which lines wrap in vanilla mode
set wrapmargin=10
" always display row/column info
set ruler
" don't highlight all search terms, just find them
set nohls
" create backup files with a tilde extension
set backupext=~
set backup
"
" Abbreviations
"
" auto-correct a favorite typo
iab teh the
"
" special key mappings
"
" use F5 to reformat paragraphs
map <F5> !}fmt <CR>
"
" Specific settings for different types of files
"
" define some specific settings for when editing C files
au FileType c setlocal shiftwidth=4 softtabstop=4 expandtab
```

As you can see, lines that begin with a single double-quotation mark enable you to intersperse comments in your .vimrc file.

NOTE **As mentioned previously, the default version of vim that is installed on Ubuntu and Kubuntu systems is a small package called `vim-tiny`. The .vimrc file given as an example uses several configuration settings that are only available if you have manually installed the vim package, which provides many enhancements to vim.**

Unfortunately, a complete discussion of the possible contents of a .vimrc file would require its own book. Some excellent examples of much more complex vim configuration files that you can use as a reference when customizing your own are available on the Web at the following URLs:

- `www.stripey.com/vim/vimrc.html`
- `www.cs.mcgill.ca/~navindra/editors/sven.vimrc`

Graphical Versions of vi

Many people, especially former users of systems such as Microsoft Windows and Mac OS X, expect applications to have a graphical user interface. This poses an interesting problem for applications such as vi and Emacs, which were originally designed for use with terminals back when a mouse was just a pesky rodent. To make people feel more at home with powerful tools such as vi (and to drag them, often kicking and screaming, into the 21st century), various graphical interfaces for vi have been developed over the years. Figure 15-6 shows gvim, which is a version of vim to which menu support has been added.

FIGURE 15-6

GNOME vim, a graphical version of vim



NOTE **GNOME vim is not installed by default on your Ubuntu system, but can easily be installed from the default repositories using the Synaptic Package Manager, as explained in Chapter 23, "Adding, Removing, and Updating Software." GNOME vim is contained in the package `vim-gnome`.**

Over the years, other graphical versions of vim have also appeared, most notably kvim, which is a version of vim tailored to the KDE desktop environment. However, kvim appears to be defunct; most of the kvim developers having moved to an open source project known as Yzis. Yzis is a new editor (inspired by vim) that is designed in a more modular fashion to more easily implement graphical extensions. For more information about this project, see `www.yzis.org`.

NOTE **If you really want to run a vi-clone in a graphical environment (not counting a Konsole or xterm window) but don't want to install all of the GNOME dependencies that GNOME vim requires, the Ubuntu repositories provide packages for elvis, a vi clone that uses X11 widgets rather than specifically requiring GNOME or KDE. It therefore works fine under KDE but doesn't bring along any GNOME baggage.**

More Information About vi and vim

Because vi is one of the oldest text editors available for UNIX-like systems, you can find numerous sources of information about it on the Web, including many that are specific to vim. Some of my favorites are the following:

- Vim Home Page: `www.vim.org/`
- Vi Lover's Home Page: `www.thomer.com/thomer/vi/vi.html`
- Introduction to Display Editing with vi (the original Bill Joy/Mark Horton paper): `http://ex-vi.sourceforge.net/viin/paper.html`
- Vim Regular Expressions: `www.geocities.com/volontir/`

A simple Internet search will turn up hundreds of other sites containing tips, tricks, and configuration settings for customizing your version of vim. However, remember that one of the primary advantages of using a vi-like editor is that it can be found on any Linux-like system, and generally works the same way on all systems. If you customize your copy of vim to death and grow to depend on your customizations, you might as well be using Emacs.

Using Emacs

The history of Emacs was described earlier in this chapter — now, let's meet our new best friend. As I mentioned earlier, I am a long-time Emacs user, and find it to be "the right editor" for me. Your mileage may vary, but no one can deny the power and flexibility of Emacs. Many of the old arguments against using Emacs, such as its memory requirements, simply don't matter anymore thanks to the generally increasing memory requirements of desktop systems and improvements in the type of hardware that most personal computer systems have today. As far as I'm concerned, the simple fact that Emacs is a modeless editor that starts up in edit mode, enabling you to create files without having to do anything fancy, is a great argument for using Emacs.

TIP **For some unknown reason, the Emacs package is not installed by default on Ubuntu systems. You will have to install it using any of the software update tools discussed in Chapter 23, "Adding, Removing, and Updating Software," before you can experiment and become one with Emacs. The Emacs package is a meta-package that always depends on the latest Emacs packages, which were based on Emacs version 22 when this book was updated. Installing Emacs is the first thing that I do after installing any Ubuntu/Kubuntu system. From my point of view, a Linux system without Emacs is a contradiction in terms.**

By default, the version of Emacs that is available for your Ubuntu system is integrated into the X Window system, and therefore automatically starts up its own window, which provides both an extensive set of menus and a graphical toolbar to simplify selecting frequently used commands. Figure 15-7 shows an Emacs window, displaying a text file that represents the beginning of this chapter.

The Emacs editor can also be started in non-graphical mode, which can be handy if you simply want to do some quick editing, or if you've connected to a remote machine and don't want to enable the remote system to create windows on your screen. To start Emacs in non-windowed mode, use the `-nw` command-line option when starting Emacs. Figure 15-8 shows Emacs editing the same file as that shown in Figure 15-7, but in the context of an xterm window.

FIGURE 15-7

Emacs on an Ubuntu system

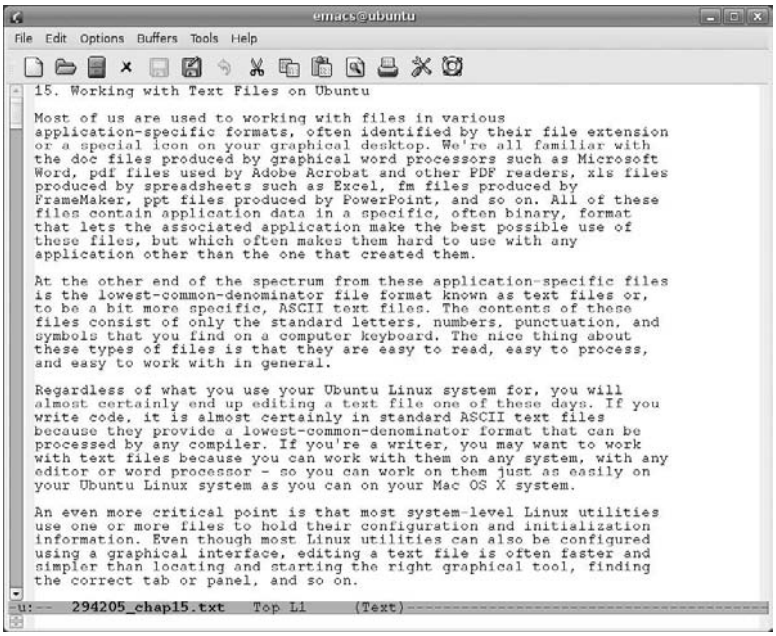
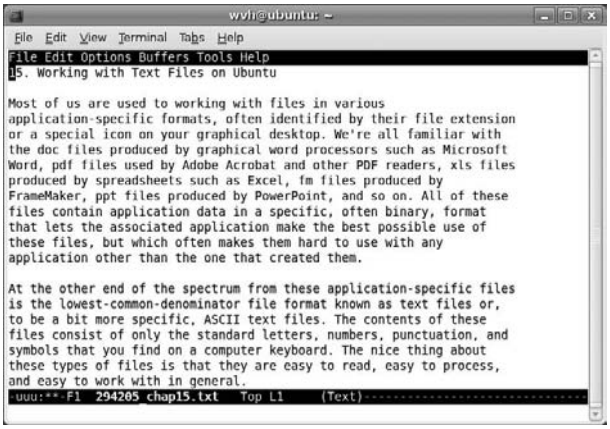


FIGURE 15-8

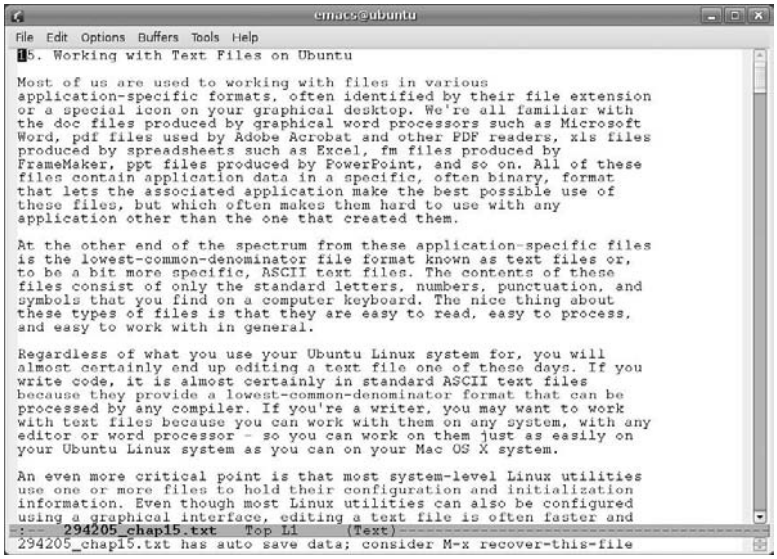
Emacs in an xterm window on an Ubuntu system



One of the best features of Emacs is its customizability, which extends to its graphical display. As an example, Figure 15-9 shows Emacs with the graphical customizations that I normally use, with the toolbar hidden and the scroll bar moved to the right of the Emacs window (as you can see, it's on the left in Figure 15-7).

FIGURE 15-9

Emacs showing graphical customization

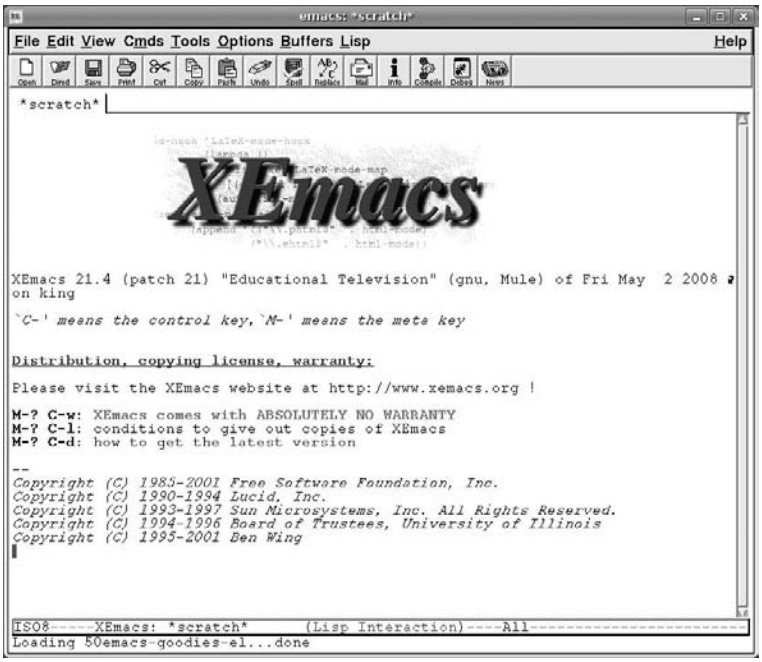


If you're a fan of graphical customization and are interested in being able to use a variety of fonts in your Emacs session, try a version of Emacs known as XEmacs. Figure 15-10 shows the default XEmacs startup window, which highlights the fact that XEmacs provides great support for multiple fonts, provides much more extensive toolbar and icon support, and even supports displaying graphical images within an XEmacs window.

XEmacs was originally the work of the folks at Lucid, Inc., but is now an open source project that is hosted at `www.xemacs.org`. XEmacs was originally created to provide a more graphically — and mouse-aware — version of Emacs, primarily because the first release of GNU emacs to intrinsically support the mouse in a meaningful way, Emacs 19, took forever to be released. Because of several legal issues that would numb my brain to discuss here, the XEmacs code has never been merged back into the official GNU emacs, and it exists as a separate entity today. Regardless, XEmacs is a fantastic version of Emacs that many people find much more attractive and usable than GNU emacs, and is therefore well worth looking into. XEmacs is not officially supported by the Ubuntu folks but is provided in the multiverse repository. The primary XEmacs package at the time that this book was written is named xemacs21. You can install Xemacs using any of the mechanisms discussed in Chapter 23, "Adding, Removing, and Updating Software."

FIGURE 15-10

The startup screen for XEmacs, the "other" Emacs



A Few Words About Emacs Commands

Before I tell you how to do anything in Emacs, it's important to have a quick discussion of Emacs commands, how they work, and how they're represented in this book.

Because Emacs is a modeless editor, typing any standard keyboard character actually inserts that character into the file that you're working on. That makes perfect sense in a graphical environment where you can use menu commands to perform most Emacs commands (as long as they're listed on a menu). However, because Emacs predates the use of graphical environments, and therefore menus, there has to be another way to issue basic Emacs commands.

All available Emacs commands are bound to a combination of keystrokes that use the modifier keys on your keyboard. These modifiers are keys such as Ctrl (known as Control), Esc (known as Escape), and the Alt and Shift keys. If you think about how a computer works, you'll realize that these keys, when combined with the standard alphanumeric and symbol keys, change the signal that is sent to your system when one of these standard keys is sent. Using the Ctrl, Escape, and Alt keys does exactly the same sort of thing as what happens when you press the Shift key when typing a letter or number key on your keyboard — it sends a different signal than what that key normally sends.

Emacs commands are therefore bound to key combinations that are represented as things like Ctrl+S, Esc+S, Alt+S, Ctrl+Shift+S, and so on. This is one of the things that the users of editors like to make fun of because watching someone type in Emacs can be quite amusing if they are executing a continuous

mixture of text insertion keystrokes and Emacs commands. As an example of this, one of my old office mates described my typing as looking like an "epileptic spider." Of course, the fact that I don't touch-type didn't help me any, but his point was still both amusing and well taken, and I am clearly still amused (but have not changed). This has something to do with how different keyboard modifier keys work.

The Ctrl key works as a direct modifier by being held down while another key is pressed. This is usually represented as Ctrl+S, which means "hold down the Ctrl key, press the letter s, and then release the Ctrl key." The Ctrl key is much like the Shift key in this respect, except that there is no equivalent "Ctrl Lock" key to confuse and irritate people.



A printed sequence beginning with the caret symbol (^) is often used as a shortcut when referring to the Ctrl key. For example, ^s and Ctrl+s are two ways of representing the same thing. This simplifies having to continually type Ctrl+S, and also prevents anyone from being confused about whether to type the + (you don't). Some books also use Ctrl+s, Ctrl-s, or simply C-s as other shortcuts for Ctrl+S. I'll stick with Ctrl+S or, occasionally, ^s. Emacs can be confusing enough without drowning in a sea of duplicate representations.

On the other hand, the Escape and Alt keys work as prefix keys. To use these keys, you press and release them once, and then press another letter. These are still represented as Esc+S and Alt+S, but the keyboard interaction is very different. Commands associated with the Escape key are therefore often referred to as an *Escape sequence* to reinforce the notion that these are a sequence of key presses, not a simultaneous combination of multiple keys.

This would be pretty simple if Emacs commands were limited to a single Ctrl command, Escape sequence, and so on. However, this would limit Emacs to having a hundred or so possible commands, which is hardly sufficient. Therefore, most Emacs commands are combinations of Ctrl keys, Escape sequences, Alt sequences, and the Shift key. This gives you an essentially infinite number of available commands, which can be hard to remember, leads to derogatory comments like those of my old office mate, and amuses everyone with potential key sequence descriptions like Esc+S+Ctrl+Shift+S. However, I feel that this is a small price to pay for the power of a zillion possible commands.



One thing that confuses users who are familiar with text editors or word processors on other systems is that Emacs does not follow the traditional conventions for things such as cutting, copying, and pasting text. This is unfortunate, but the reason for this is because Emacs predates these conventions, and those key sequences are already used for other things in Emacs. Who would you rather tick off — thousands of Emacs users who are potentially unruly in the first place, or new users who are new to Emacs anyway? I thought so . . .

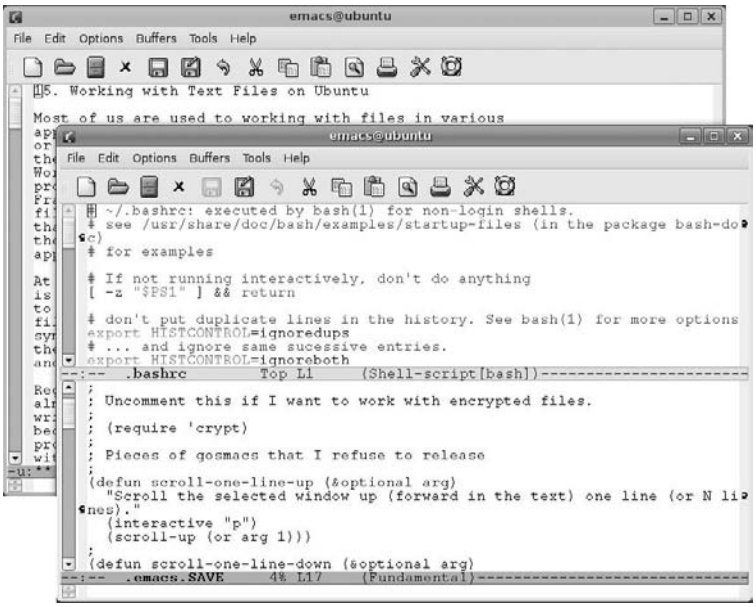
Emacs Terminology

Because GNU emacs is designed to interact with a graphical window manager, I have to use slightly different terminology to describe what you see on the screen. This section explains that terminology so that I can use it without guilt throughout the rest of this chapter.

Figure 15-11 shows two windows on an Ubuntu system, each running Emacs but displaying different files. Each Emacs window has its own menu bar and toolbar. The window in the foreground displays two different files, each in its own buffer. Each Emacs buffer has a mode line at the bottom, which provides information about the file being viewed in that buffer and its state. At the bottom of each Emacs window is a status line in which you interact with Emacs, or in which Emacs displays any messages displayed by a particular Emacs function.

FIGURE 15-11

Windowing terminology for GNU emacs



NOTE **The term "window" has a different meaning in this section than the way in which I used it in the portions of this chapter that discussed vi and vim. This is because Emacs is natively aware of the X Window system and therefore can display multiple windows on the screen, while vi can only divide a terminal application screen into multiple sections. In the context of Emacs, a vi window is equivalent to an Emacs buffer.**

Starting and Exiting Emacs

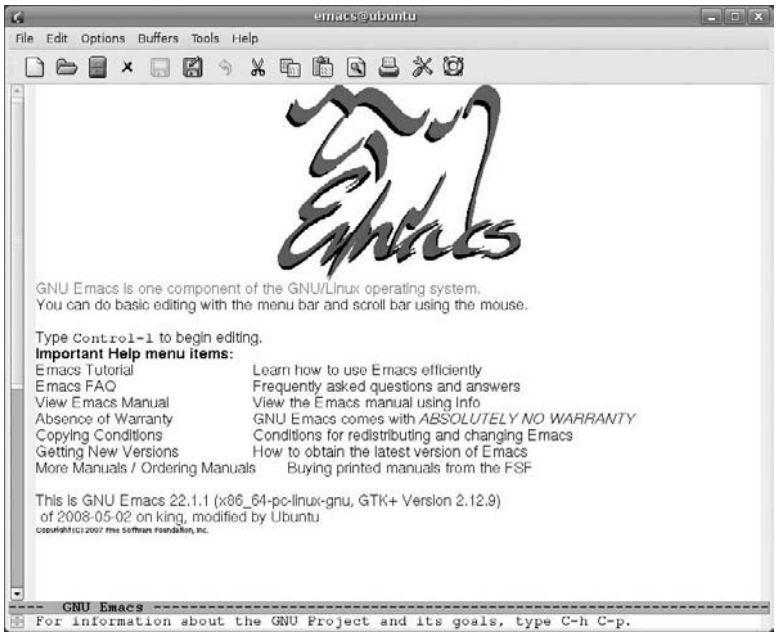
As with vi, to begin editing a file in Emacs, you can do either of the following:

- Select the Emacs 22 (GTK) command from the Applications ➪ Accessories menu.
- In any xterm or GNOME Terminal window, simply type the command Emacs filename, where filename is the name of the file that you want to edit. By default, starting Emacs from a terminal window on an Ubuntu system creates a new window for the Emacs application. As a result, you will probably want to follow the name of the file that you want to edit with an ampersand to start your Emacs session in the background and return control to the terminal window so that you can type additional commands there while Emacs is running.

In either case, Emacs will create an initial window displaying any file that you are editing, as shown in Figure 15-7. If you start Emacs from the GNOME menus, Emacs displays an initial welcome screen, as shown in Figure 15-12.

FIGURE 15-12

The startup screen for GNU emacs



To exit Emacs, you can use either the Exit Emacs command on the File menu the ^x^c (as a refresher, this means to hold down the Ctrl key and press the letter x, then press the letter c, and then release the Ctrl key) key combination.

If you have not modified any files while working with Emacs or have already saved all of the changes that you have made, Emacs exits immediately. However, if you attempt to exit from Emacs using menu commands and have not saved any of the changes to the file(s) that you've been editing, Emacs displays the confirmation dialog shown in Figure 15-13 for each file that contains unsaved changes.

To proceed without saving your changes, click No. To save your changes and proceed, click Yes. If you save your changes by selecting Yes, Emacs exits immediately once those changes have been saved for each modified file that you were working on. If you click No, Emacs displays yet another confirmation dialog, asking if you're really, really sure that you want to exit without saving your changes, as shown in Figure 15-14. Emacs doesn't quite beg you to save your files, but it makes every attempt to ensure that you know what you're doing and aren't accidentally exiting without saving the great American novel.

This same exchange occurs when you're attempting to exit Emacs using the equivalent command (^x^c), but it occurs in the Emacs message buffer instead of as a dialog, as shown in Figure 15-15.

FIGURE 15-13

Changed file notification when exiting Emacs graphically

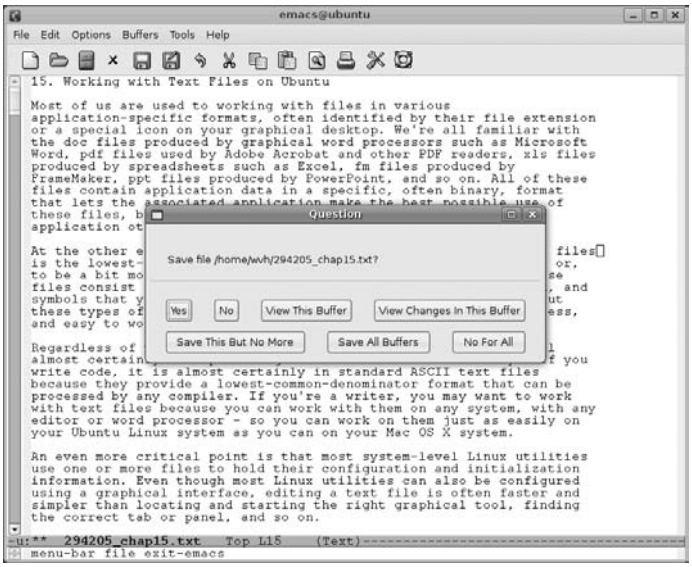


FIGURE 15-14

Final confirmation dialog when exiting Emacs graphically

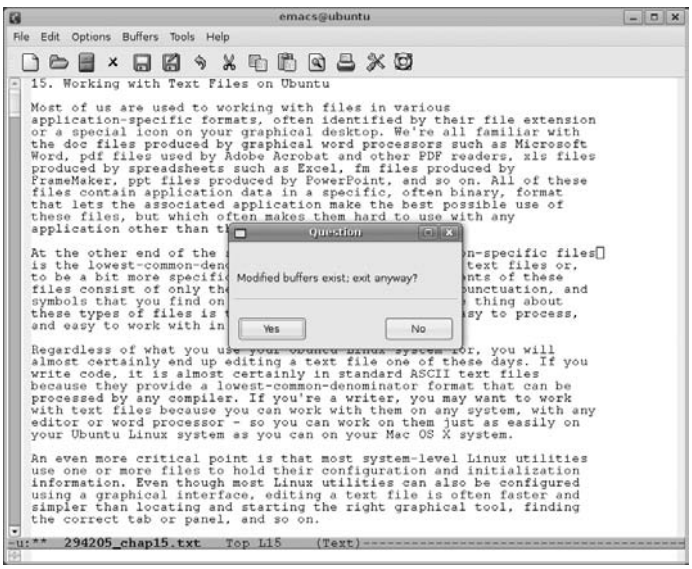
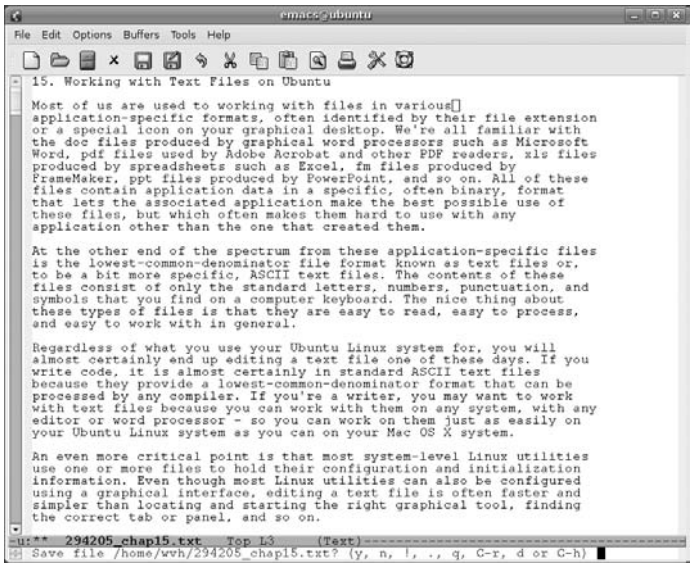


FIGURE 15-15

Final confirmation prompt when exiting Emacs using the keyboard



Moving Around in Emacs

There are many different keys for moving the insertion cursor in Emacs:

- ^a: Go to the beginning of the current line.
- ^e: Go to the end of the current line.
- ^f: Move forward one character.
- ^b: Move backward one character.
- ^p: Move to the previous line.
- ^n: Move to the next line.
- ^v: Scroll down one screen.
- Esc+b: Move backward one word.
- Esc+f: Move forward one word.
- Esc+v: Scroll up one screen.
- Esc+<: Go to the beginning of the current buffer.
- Esc+>: Go to the end of the current buffer.

These are the most common keyboard commands for moving the cursor around within the current window. When using Emacs in any graphical mode, you can also scroll the screen by clicking and dragging the button in the scroll bar, or simply move the cursor by pointing and clicking.

For information about the commands used to switch or move between multiple buffers that are displayed on the screen, see the section later in this chapter entitled "Working with Multiple Windows and Buffers in Emacs."

Cutting, Copying, and Pasting in Emacs

Because Emacs is aware of the X Window system and therefore supports moving the cursor using the mouse, there are two different ways to select the text that you want to copy or move:

- By positioning the cursor using the left mouse button, holding down the left mouse button, and dragging the cursor to select a portion of your text. If you move the cursor into the mode line of any Emacs buffer, Emacs will automatically scroll the text while continuing the selection operation.
- By setting a mark at the current cursor position by using the Ctrl+@ or Ctrl+spacebar commands, and then moving the cursor to the location that delimits the end of the text that you want to cut or copy.

Once you have identified the region that you want to copy or cut, you can copy that region to an internal paste buffer by pressing Esc+W, or cut the selected text to the paste buffer by pressing Ctrl+W. You can then move the cursor to the position where you want to paste the copied or deleted text, and press Ctrl+Y to paste the text at the new location.



If you simply want to cut the text from the current cursor position to the end of the current line, you can press ^k to kill that text, and then press ^y to paste it back in elsewhere.

Searching for and Replacing Text in Emacs

By default, the ^s key is bound to an incremental search mode that is extremely convenient to use. After you press ^s, the message I-Search: displays in the Emacs message buffer, and Emacs searches forward in the current buffer for the text that you type there, finding the first match for whatever string you are currently typing. For example, if you are searching for the string "cantaloupe", as you type each letter of that word, Emacs searches forward for the first instance of the letter c, then the first instance of the letters ca, the first instance of the letters can, and so on, much like the incremental search behavior in the Firefox web browser. You can press ^g at any time to abort the search and return to the cursor position at which the search began. Typing any other Emacs command (with the exception of ^g, that is) terminates the search and leaves the cursor at the current location in the buffer.

Working with Multiple Windows and Buffers in Emacs

If you are using Emacs in graphical mode, switching between buffers or windows using the mouse is pretty easy — you just left-click the mouse in the window or buffer that you want, and voilà! However, if you're already working in a window that's been split into multiple buffers (or if you're using Emacs in no-window mode), it's useful to know how to switch between different buffers.

To switch to another buffer that is currently displayed on your Emacs screen, press the ^x^o command. This command cycles through all of the buffers that are currently displayed on your screen.

If you want to switch to another buffer regardless of whether or not it is displayed on the screen, you can press ^x+b (Ctrl+X, and then the letter b) and enter the name of the buffer that you want to switch to at the prompt in the Emacs status line.

If you don't happen to remember the name of the other file that you were working on or are using a temporary name, you can press ^x^b to list all available buffers, and then middle-click on the name of the buffer that you want to switch to.

Customizing Emacs

If you've ever used Emacs on another user's computer, you've probably noticed that it doesn't seem to work the same way as it does on your system. The reason for this is that Emacs is the most configurable editor in existence. Not only does Emacs provide a rich configuration language for changing the commands that different keys are associated with (known as *key bindings*), but it includes a complete implementation of the Lisp programming language that makes it easy for serious Emacs users to write their own commands or modify the behavior of existing ones.

Most Linux systems store all per-user Emacs customization information in the file .emacs in the user's home directory. As mentioned throughout this chapter, Emacs' customizability and flexibility accounts for much of its popularity. The next few sections highlight the basic types of customizations that you can make in an Emacs configuration file, providing examples of each. Several excellent sites on the Web provide much more detailed information about customizing Emacs, such as `http://jeremy.zawodny.com/emacs/emacs-4.html` and `http://linuxplanet.com/linuxplanet/tutorials/3166/4/`.

Changing Key Bindings

Whenever you execute an Emacs command, you are actually executing a Lisp function. As explained earlier, Emacs is configured to execute specific commands in response to combinations of the Ctrl or Escape keys and the standard keys on your keyboard. These are known as key bindings because they associate (bind) a specific function with a specific key sequence. However, if you're already familiar with another editor that also used commands consisting of combinations of the Ctrl or Escape keys and the standard keys on your keyboard, your fingers are probably used to pressing certain keys to execute certain commands. The most common customization made to Emacs is therefore to change the keys to which commonly used commands are associated.

Although emacs enables you to change key bindings globally or within a specific mode, it is more common to customize a specific key binding so that it works regardless of the mode in which you are using Emacs. The emacs configuration command to globally set a key binding is global-set-binding. For example, to globally set the key sequence Ctrl+Z to a function that scrolls the current buffer up one line rather than attempting to suspend Emacs, you would put the following command in your ~/.emacs file:

```
(global-set-key "\C-Z" 'scroll-one-line-up)
```

In Lisp fashion, you must preface the name of the function that you are referring to with a single-quotation mark (not the back quote). If you're interested, the scroll-one-line-up function is provided as an example in the section entitled "Defining Your Own Functions" later in this chapter.

NOTE **When you are specifying key bindings in an Emacs configuration file, the Ctrl key is usually represented by `\C-` and the Escape key is usually represented by `\M-`.**

Setting Variables

Emacs uses several internal variables to control its behavior. You can modify these variables using the setq command. For example, to cause Emacs to scroll more smoothly in one-line increments, you could add the following line to your ~/.emacs file:

```
(setq scroll-step 1)
```

As another example, the default settings in Emacs cause it to automatically save your files each time you have typed or modified 300 characters. To decrease this value to every 100 changes, you could add the following line to your ~/.emacs file:

```
(setq auto-save-interval 100)
```

Specifying Modes

As discussed earlier, Emacs provides different key bindings and functions based on the types of files that you are using, known as *modes*. The default mode used by Emacs when you start Emacs without specifying the file that you want to edit is known as *fundamental mode*. You may want to customize Emacs to always use a different mode by default, regardless of the name of the file. For example, to make text mode the default mode each time you start Emacs, you could add the following line to your ~/.emacs file:

```
(setq default-major-mode 'text-mode)
```

As you can see from this example, the name of the default mode in Emacs is defined through an Emacs variable, default-major-mode, which in this case requires the name of an Emacs function rather than a numeric value as shown in the examples in the previous section.

Each Emacs mode also enables you to specify other actions whenever you use that mode. For example, to turn on auto-fill (wrapping words to the next line when you reach a certain column on your screen), you could add the following line to your ~/.emacs file:

```
(add-hook 'text-mode-hook 'turn-on-auto-fill)
```

Changing File Associations

As discussed earlier, Emacs automatically enters specific modes when you open files with a specific extension, just as many GUIs do when you click on a specific type of file. You may occasionally want to have Emacs automatically enter a specific mode when you open files that the specified mode is not traditionally associated with. For example, Emacs automatically enters text mode when you edit files with the .txt or .text extensions. But what if you typically name text files with the .inf extension and want Emacs to automatically enter text mode when entering files of that type?

The file extensions associated with different Emacs modes are stored in a list consisting of extension/mode-name pairs. Therefore, to add the .inf extension to the list of extensions associated with text mode, you would simply append that extension to the list of extensions associated with text mode, as in the following example:

```
(setq auto-mode-alist
  (cons '(".inf" . text-mode) auto-mode-alist))
```

This example uses the Lisp cons function to append the pair .inf and `text-mode` to the list of automatic file associations, and then uses the setq function discussed earlier in this chapter to set the `auto-mode-alist` variable to the new, expanded list.

Defining Your Own Functions

The fact that Emacs includes an internal implementation of Lisp makes it relatively easy to define your own functions and provides a practical reason for learning Lisp if you've always wanted one. Lisp functions are defined using the `defun` command, which takes a list of five values as arguments. Without turning this into a Lisp tutorial, let's just quickly look at a sample function definition and examine its components.

The following Emacs Lisp code defines a simple function that scrolls the contents of the current window up some number of lines, with the default number being 1:

```
(defun scroll-one-line-up
  (&optional arg)
  "Scroll the selected window up one or N lines."
  (interactive "p")
  (scroll-up (or arg 1))
)
```

The name of the function that you are defining is the first argument to the `defun` command. The second argument is the list of arguments that are used by the function that you are defining. In this case, this list uses the `&optional` argument to state that this function can take one argument, but that the argument is optional. The third argument to the defun command is an informational message about the function, enclosed in double quotes. The fourth argument defines whether the function is interactive (p means true in Lisp). The fifth argument is the actual Lisp code for the function that you are defining, which in this case executes the built-in Emacs scroll-up function a single time or the number of times you specified as a numeric argument to the function.

This section won't exactly make you a Lisp wizard, but just knowing that Emacs includes a complete Lisp implementation can be quite empowering. As you can see from this example, defining simple functions is not all that complex. The power and flexibility of Emacs provides you with the chance to add that function you've always wished for in a text editor!

Getting More Information About Emacs

The Emacs editor contains its own tutorials and help files: Type `M+x help` to begin. These include a "learning by doing" tutorial. There are also plenty of Emacs tutorials out there, some of which are written from the beginner's point of view. The official GNU Emacs manual is available from `www.gnu.org/software/emacs/manual/`. It can also be purchased in book form. There is also an Emacs Wiki at `www.emacswiki.org`.

Using gedit

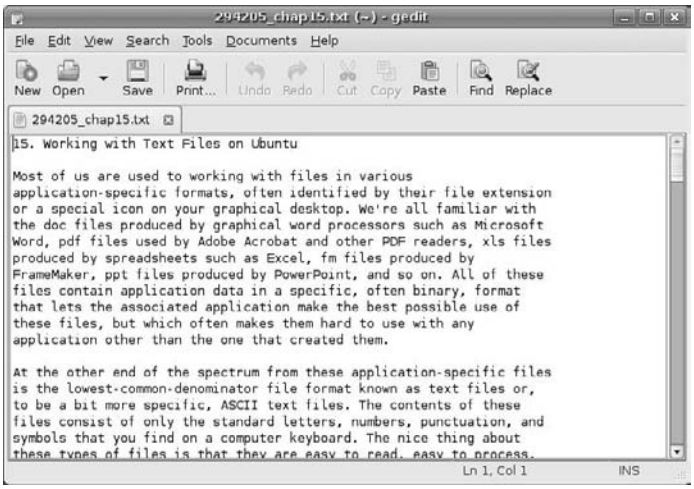
Now that I've discussed the favorite two classic UNIX/Linux text editors and their attendant philosophies, it's worth mentioning that Ubuntu also provides text editors that you can use without a class in UNIX history and pledging your undying allegiance to Richard Stallman or Bill Joy. Both the GNOME and KDE desktops provide easy-to-use graphical text editors called gedit and kedit, respectively. Figure 15-16 shows gedit displaying the same file used to illustrate the vi and Emacs editors earlier in this chapter, namely, the text of this chapter as I was writing it.

The gedit editor is completely mouse- and menu-driven, and follows the standard keyboard conventions for most graphical editors across platforms. These include the following:

- Ctrl+C: Copy selected text.
- Ctrl+N: Open a new file.
- Ctrl+S: Save the current file.
- Ctrl+V: Paste copied or cut text.
- Ctrl+X: Cut selected text.

FIGURE 15-16

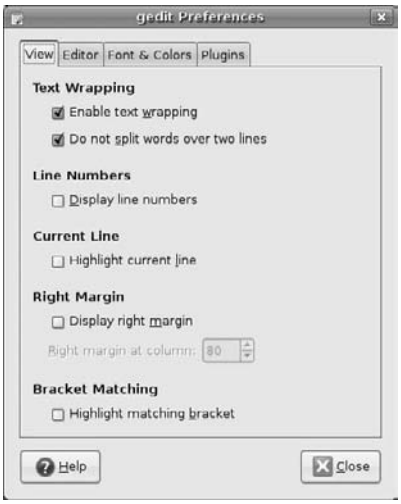
The gedit text editor



As a graphical GNOME application, gedit provides convenient and more traditional dialogs for interacting with, moving around in, and selecting files and directories. It can also be customized to suit your specific editing needs by selecting the Edit menu's Preferences command, which displays the dialog shown in Figure 15-17.

FIGURE 15-17

gedit's Preferences dialog



The tabs on this dialog provide access to different sets of configuration options, so that you can make gedit work the way that you want it to. For example, word-wrapping is great for documents like this one, but you probably don't want to use it when you are writing programming code. The most interesting of the preferences tabs is the Plugins tab, which enables you to activate or deactivate different features that are provided as extensions to the basic capabilities of gedit.

As with most GNOME applications, extensive online help for gedit is available by selecting the Help menu's Contents command.

Popular KDE Text Editors

If you are using Kubuntu, you can still use the Emacs or vi editors that were discussed earlier in this chapter, but gedit isn't available on Kubuntu systems unless you install it manually. I strongly recommend against doing that, for two reasons. First, installing all of the GNOME dependencies just to run a simple text editor such as gedit is like buying a supercomputer to play pong. More important, Kubuntu already offers several graphical text editors that are designed for KDE and should easily satisfy your graphical text editing needs.

All of the KDE text editors discussed in this section are completely mouse- and menu-driven, and follow the standard keyboard conventions for most graphical editors across platforms. These standard keyboard conventions include the following:

- Ctrl+C: Copy selected text
- Ctrl+N: Open a new file
- Ctrl+O: Open an existing file
- Ctrl+P: Print the current file
- Ctrl+S: Save the current file
- Ctrl+V: Paste copied or cut text
- Ctrl+X: Cut selected text

As noted, these keyboard conventions are the same as those in popular graphical word processing packages for Microsoft Windows and Mac OS X systems (although on the latter, the Command key is used in place of the Ctrl key), which simplifies moving to a Kubuntu system because your fingers don't have to learn a whole new set of keyboard shortcuts.

Using KEdit

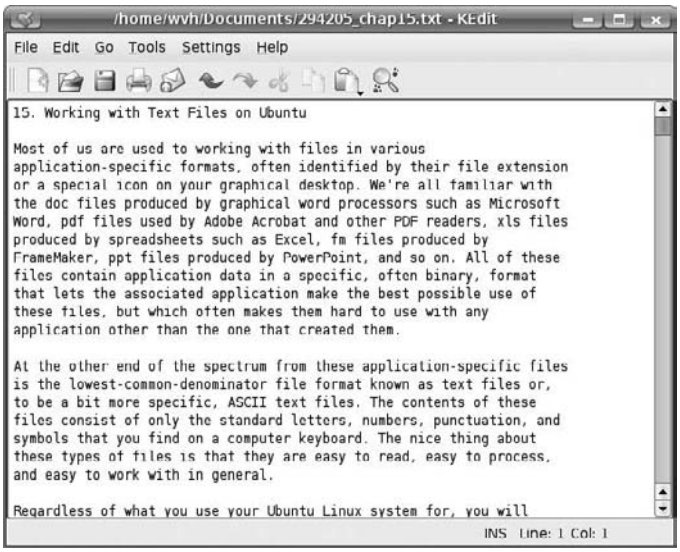
KEdit is a simple text editor for Kubuntu that is analogous to gedit — a fast, minimal text editor that is easy on system resources. Figure 15-18 shows KEdit displaying the contents of this chapter as I was working on it.

KEdit is no longer installed by default on KDE 3-based Kubuntu systems, but can be added by installing the kedit package using the Adept Manager. (Kedit doesn't seem to be available for KDE 4-based Kubuntu systems.) This is not a criticism of Kedit, but simply acceptance of the fact that it doesn't offer significant advantages over KWrite in terms of speed, and KWrite provides features that KEdit does not. Installing KEdit does not automatically create a menu item for the application, so you must start it by:

- Pressing Alt+F2 and entering **/usr/bin/kedit** as the command to execute
- Entering /usr/bin/kedit in a Konsole or xterm window
- Creating a menu or panel launcher for /usr/bin/kedit and then selecting that item

FIGURE 15-18

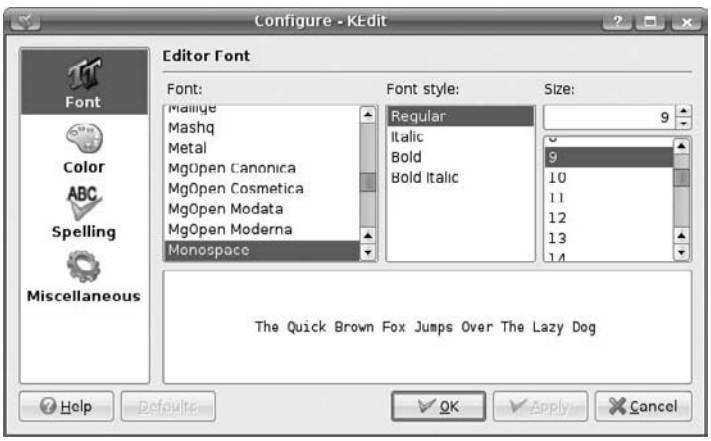
Editing a text file in KEdit



Although designed to work out of the box with no customization required, some aspects of KEdit can be configured using the dialog displayed by the Settings menu's Configure KEdit command, shown in Figure 15-19.

FIGURE 15-19

Configuring KEdit



The KEdit configuration dialog lets you configure things such as the fonts used in the edit window, the foreground and background colors, the spell-checking utility and encoding used by for spellchecking. It also enables you to configure miscellaneous settings such as how word-wrapping works and whether KEdit automatically creates backups when you save updates to a file.

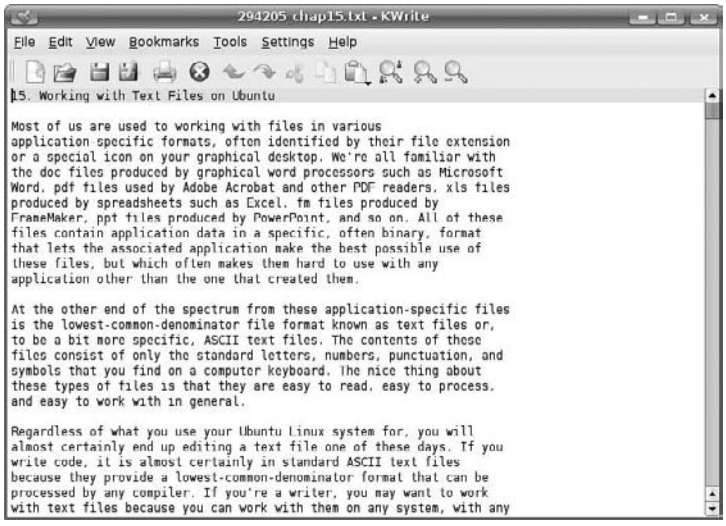
Extensive help for KEdit is available by pressing F1 in a KEdit window, or by selecting the KEdit Handbook menu command from KEdit's Help menu.

Using KWrite

KWrite is a small, fast text editor for Kubuntu that is also extremely customizable. On Kubuntu systems running KDE3.*x* systems, KWrite is provided in the same Kubuntu package as the more complex and powerful Kate editor (discussed in the next section), and is therefore installed by default on Kubuntu systems. (KWrite is also installed by default on Kubuntu systems running KDE 4.*x*, but is a standalone package there.) Figure 15-20 shows KWrite displaying the contents of this chapter as I was working on it.

FIGURE 15-20

Editing a text file in KWrite



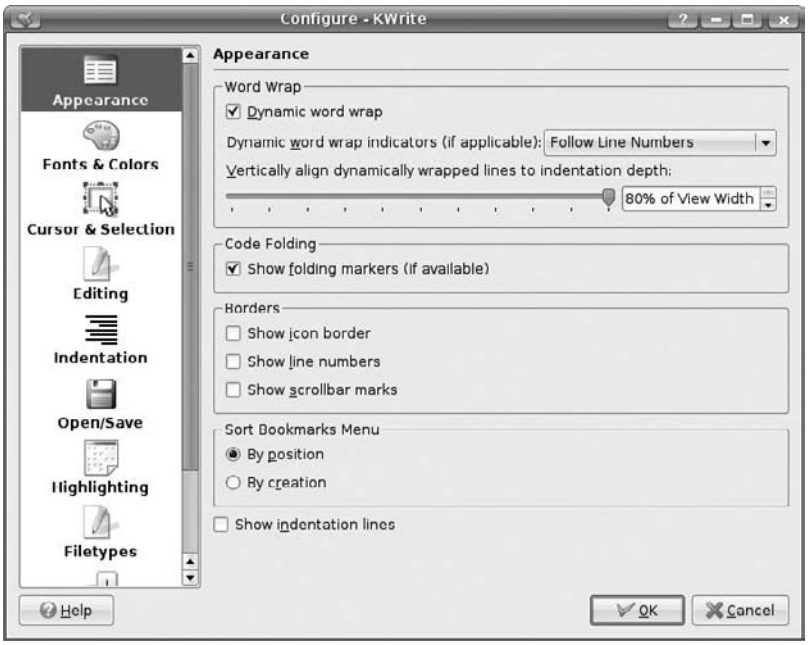
KWrite is not just a text editor for text such as this book. KWrite provides built-in support for writing code, including useful features such as syntax highlighting.

As with KEdit, KWrite starts quickly and is quite easy on system resources. Unlike KEdit, you can configure a vast number of aspects of KWrite. Figure 15-21 shows the KWrite Configuration dialog, which you can display by selecting the Settings menu's Configure Editor menu command.

In addition to enabling you to customize its current configuration, KWrite also enables you to extend its core functionality by adding plug-ins, which can then be enabled by selecting the Plugins icon by scrolling down to the bottom of the toolbar shown at the left of Figure 15-21.

FIGURE 15-21

Configuring KWrite



Extensive help for KWrite is available by pressing F1 in a KWrite window, or by selecting the KWrite Handbook menu command from KWrite's Help menu.

Using Kate

Kate, the KDE Advanced Text Editor, is exactly that: a graphical text editor that is highly configurable and provides advanced capabilities such as integrated support for writing code in a variety of programming languages via syntax highlighting, editing multiple files at one time, splitting the editing window horizontally or vertically to enable you to easily work in multiple portions of the same file at one time, and much more. Figure 15-22 shows Kate displaying the contents of this chapter as I was working on it.

Figure 15-22 shows the primary feature of Kate: a left pane that lists the files that are currently open, with the main pane being the editing window for the currently selected file.

You can start Kate in KDE 3-based Kubuntu systems by selecting the K ➪ Utilities ➪ Kate — KDE Advanced text Editor menu command. On KDE 4-based Kubuntu systems, you select the K ➪ Applications ➪ Kate Editor menu command, or from the command line if the menu item is not available.

Figure 15-23 shows Kate's Configuration dialog, which you can display by selecting the Settings menu's Configure Kate menu command.

FIGURE 15-22

Editing text files in Kate

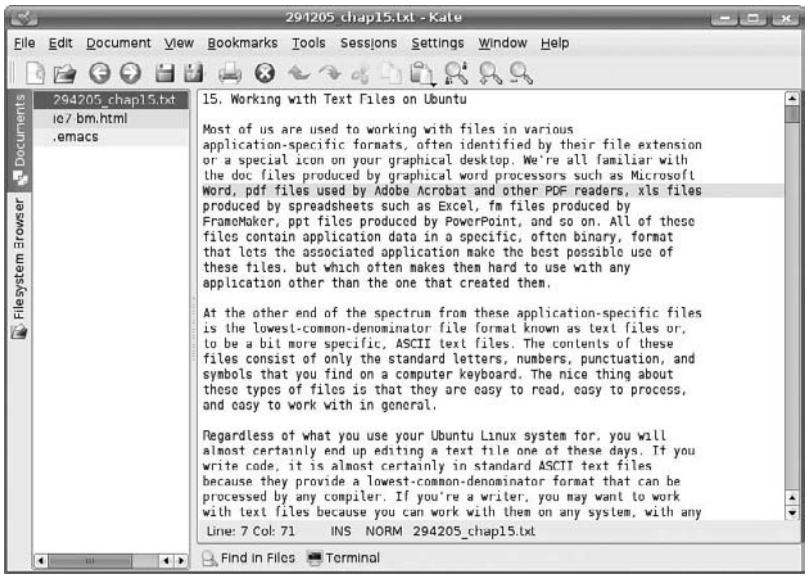
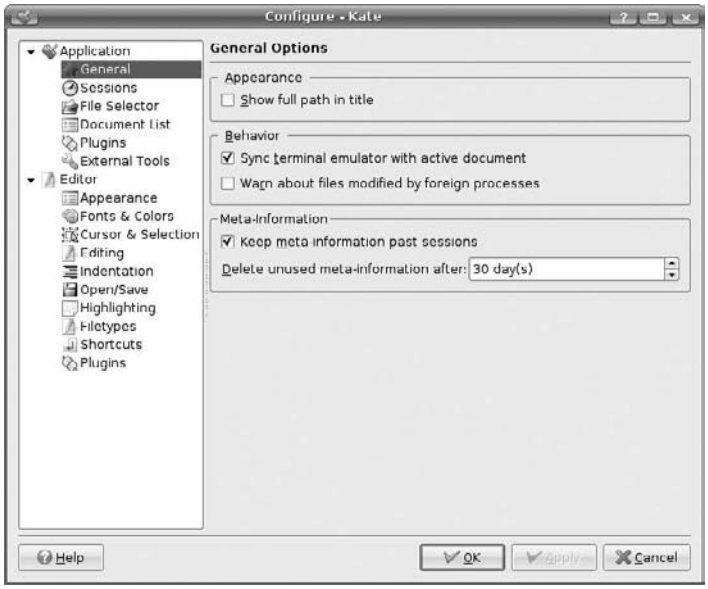


FIGURE 15-23

Configuring Kate



As you can see from this figure, Kate provides substantially more configuration items than KWrite, a fact that generally endears it to developers who want their editor to work a certain way.

Some interesting Kate capabilities are as follows:

- The ability to bookmark directories and easily jump to them, which can be very handy when you are working on large projects composed of many files that are scattered across multiple directories.
- Session support that enables you to quickly return to a previous editing session that reopens the file(s) that were being edited and relocated the cursor(s) at their last position.
- An integrated filesystem browser that appears in the left pane when you select it from the far left border of a Kate editing window. When trying to locate specific files, this provides a much better user experience than the traditional filesystem navigation dialogs provide by File ➪ Open commands. This is especially true if you are working on multiple files that are in widely different filesystem locations.
- An integrated command-line interface (referred to as a "Terminal") that you can expose by selecting the Window ➪ Tool View ➪ Show Terminal menu command. Personally, I find this feature more interesting than useful, but it can certainly be handy if you need to execute commands and insert their input into the file that you are editing. For example, you might find built-in terminal access to be very convenient if you were writing a book on Linux.

Extensive help for Kate is available by pressing F1 in a Kate window, or by selecting the Kate Handbook menu command from Kate's Help menu.

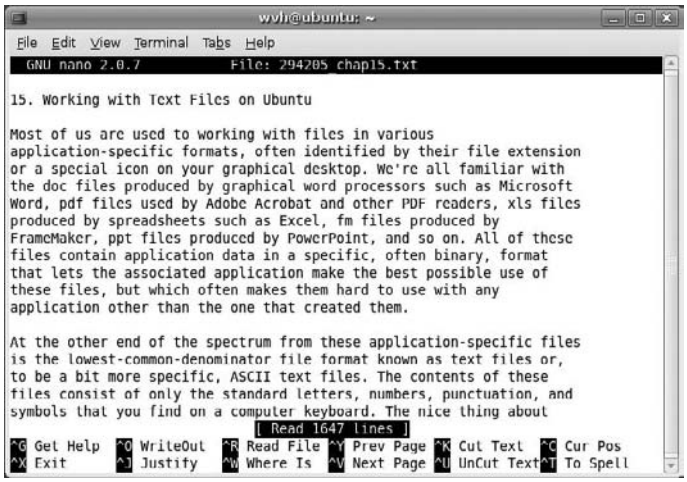
Other Text Editors for Linux Systems

Several other editors are available for Ubuntu/Kubuntu systems. By default, the following other editors are installed as part of a standard Ubuntu or Kubuntu installation:

- ed: A simple, line-oriented editor inspired by the original UNIX text editor. Be careful when you start this, because it doesn't actually display a prompt. (The command to exit is the letter q.)
- ex: Invokes the vim editor in line-oriented mode. When using ex, you can instantly switch to vim's full-screen mode by typing the keyword visual.
- nano: An enhanced clone of pico that is available under the GPL (unlike pico). One of the nicest things about nano and pico is that the bottom of any window always provides a quick summary of the most useful commands in the current context. nano (and pico) is completely unaware of the X Window system, and must be executed within an xterm or GNOME Terminal window. Figure 15-25 shows nano editing the text file for this chapter.
- pico: A simple, screen-oriented text editor that was originally introduced by the Pine mail system on UNIX systems. On Ubuntu Linux systems, this is actually just a symbolic link to the nano editor, and thus looks exactly like the screen shown in Figure 15-24.

FIGURE 15-24

The nano text editor



If you are new to Linux, you might be interested in two other popular editors that are not provided with Ubuntu but which you can compile or execute on your Ubuntu system:

- jedit: A popular, cross-platform editor for programmers. For more information, see `www.jedit.org/index.php`.
- THE: The Hessling Editor, a clone of the popular IBM mainframe editor XEDIT, THE includes support for various Linux versions of the associated REXX scripting language. See `www.lightlink.com/hessling` or `http://hessling-editor.sourceforge.net` for more information.

If you're looking for a Linux clone of a specific editor, check the Web page at `http://texteditors.org`. This site lists many available editors for different operating system (including Linux), and may help you find the exact editor that best meets your needs.

Summary

Your Ubuntu Linux system provides a variety of text editors that different types of users can quickly get comfortable with. If you are familiar with Linux or other UNIX-like systems, Ubuntu installs vim (a vi clone), and GNU emacs is easily installed using the software update commands discussed in Chapter 23, "Adding, Removing, and Updating Software." This chapter discussed in detail how to use both vi and Emacs, how to accomplish the most common tasks you'll need to do in an editor, as well as how to customize each to suit your preferences. It then provided an introduction to gedit, a fast graphical editor that is installed as part of the GNOME desktop environment. The chapter concluded with information about other editors that are available on your Ubuntu system, and provided suggestions for finding other editors that may be more to your liking.

Chapter 16

Creating and Publishing Documents

Chapter 15 discussed using text editors on Linux systems. Text editors are fine for creating documents if the document you're creating is a shopping list, To-Do list, or ransom note. For more sophisticated documents, most people nowadays expect an actual word processor, which is an application that enables you to create documents that use multiple fonts, and different styles within a font (bold, italic, and so on); create bulleted and numbered lists automatically; and provide some port for page layout using advanced constructs such as tables, multiple columns, and so on.

This chapter explains how to do various types of word processing on your Ubuntu or Kubuntu Linux system. It begins by discussing the various markup languages, which are text-format ways of creating formatted documents in a text editor on your Linux system. Although somewhat old school, markup languages are still popular, and Ubuntu supports popular, open source, markup-oriented document production systems such as TeX, and even provides a clone of the venerable troff system from bygone UNIX days.

NOTE **If you're reading this chapter and thinking, "Where the heck does he talk about markup languages that I actually care about, like HTML and XML?" I'm sorry to disappoint you. These markup languages are not discussed in detail in this book because they are more focused on Web site creation than document creation, and because the world is already knee-deep in books dedicated to these topics.**

Next, this chapter discusses more familiar types of graphical word processors, focusing on OpenOffice.org Writer, which is the word processing component of the powerful OpenOffice.org suite of applications that is featured throughout this book. OpenOffice.org is a fantastic suite of office software, and this chapter's discussion of using OpenOffice.org Writer should make even the most dyed-in-the-wool Microsoft Word, Corel WordPerfect, or Mac OS X Nisus Writer user feel right at home. This chapter also discusses KWord, a powerful combination of a word processor and a desktop publishing tool that is available

IN THIS CHAPTER

Markup Languages: TeX and LaTeX

Using OpenOffice.org Writer

Using KWord

Using Scribus

Other word processors for Linux

only if you are using Kubuntu or have manually installed KWord (and KDE dependencies) on a standard GNOME-based Ubuntu system.

Finally, this chapter discusses desktop publishing on Ubuntu, highlighting the excellent Scribus application. When you have to do serious, hard-core layout in your documents, a word processor is no longer the right tool because it doesn't provide sufficiently granular control of page elements. Desktop publishing software makes it easy to create documents that contain many different page elements, routing text from one to the other as needed, and easily inserting and manipulating graphics to create documents that look exactly the way that you need them to. Scribus is a relatively young but incredibly impressive application. Adobe InDesign, Adobe PageMaker, and Quark XPress users, hold on to your hats, but don't worry about your wallets—Scribus is open source!

Using Document Markup Languages

Long ago, when people just used terminals, word processing meant creating documents in a text editor and embedding special commands to tell some other application how to format your document for a specific target printer. The format in which these documents were created is known as a *markup language*, of which the most common example is HTML. Of course, HTML is designed to mark up documents that are targeted for use as Web pages, but the same principles apply. Creating documents in markup languages is still popular in academic and publishing circles because many institutions and journals use text files in a markup language as the lowest common denominator for submissions. This enables them to combine documents from multiple sources into a single, stylistically consistent publication without worrying about what version of Word everyone used, what fonts they'd used, and so on.

TeX (pronounced *tech*) is a typesetting program and associated markup language that is designed to produce extremely high-quality output, especially for scientific and mathematical notation. The eminent computer scientist Donald Knuth wrote TeX in 1977, largely because he needed a suitably high-quality tool for producing the first three volumes of his "Art of Computer Programming" series. As if it weren't cool enough to write your own document formatting system, Knuth originally wrote TeX in a system of his own called WEB, which is a mixture of documentation and Pascal source code in a single source file. The source code is extracted from the WEB document using a program called tangle, and the formatted source code and integrated documentation is extracted (in TeX format, of course) using a program called Weave. None of this matters if you just want to use TeX, but it is pretty entertaining and sounds like a recursive episode of "Computer Scientists Gone Wild" (available soon on DVD), so I thought you might enjoy adding this tidbit to your collection of interesting but non-essential knowledge.

Many different implementations of TeX are available today, all based on the version extracted from the original `tex.web` document, many converted into other languages, and all enhanced with the addition of extra capabilities. The source code for TeX is freely available, but in order for anything based on the TeX source code to call itself TeX, it must pass a TeX formatting torture test known as TRIP. The version of TeX provided in the Ubuntu repository (as on most Linux systems nowadays) is TeX Live, a complete TeX distribution for UNIX-like and Windows systems that is collaboratively developed by a number of TeX user groups. The home page for TeX Live is `www.tug.org/texlive`.

Installing TeX and LaTeX

Because TeX isn't for everybody, it isn't installed by default on Ubuntu or Kubuntu Linux systems. The Ubuntu repositories provide a huge number of TeX-related files, but the only one that you really need to install is the `texlive` package, which depends on many others and therefore guarantees that you'll get

everything you need for a usable TeX formatting environment. You may also want to install the appropriate `texlive-doc` package for your national language which, for English-speakers like me, would be the package `texlive-doc-en`. The `texlive-doc` packages provide FAQs, HOWTOs, tutorials, and many other documents about the TeX system.

To install TeX on an Ubuntu or Kubuntu system, use your favorite package management software, as discussed in Chapter 23, to install the texlive package. If you're interested in the documentation (and you should be, unless you're already a hard-core TeX-type), you should also install the appropriate `texlive-doc` package for your national language. Once the installation completes, exit the package manager, and you're ready to begin experimenting with TeX.

See Chapter 23, "Adding, Removing, and Updating Software," for more detailed information on installing applications using the package management software provided on Ubuntu and Kubuntu systems.

Using TeX Markup

TeX markup is quite simple. For example, the source code for a TeX document that produces formatted output that displays the string "Hello, World!" is as follows:

```
Hello, World!
\bye
```

To format this document, you simply execute the `tex` command from an Ubuntu command line, supplying the name of the file as an argument. You'll see something like the following on your screen:

```
$ tex hello_tex.tex
This is TeX, Version 3.141592 (Web2C 7.5.6)
(./hello_tex.tex [1] )
Output written on hello_tex.dvi (1 page, 228 bytes).
Transcript written on hello_tex.log.
```

Processing a file using TeX produces an output file in DVI (Device Independent) format. DVI output is an intermediate binary form that contains character positioning and font information in addition to the content of your document, and must be post-processed by another program that understands how to display this file on your screen or produce output in the specific format or printer control language used by your printer.

As you can see from the sample TeX markup for the "Hello, World!" example, you don't need many TeX commands to produce simple documents. However, the key word here is simple. To produce a slightly more complex document using different fonts, centering, and different sections, your native TeX code would look something like the following:

```
\centerline{This is a Title}
\smallskip

\centerline{Bill von Hagen}
\centerline{\tt 17 May, 2008}

\bigskip

{\narrower\noindent
This example shows some of the complexity of executing raw \TeX\
commands and is designed to illustrate how much simpler this same text
```

```
would be if I were using a higher-level macro package of some
sort. Wouldn't it be nice...\par}

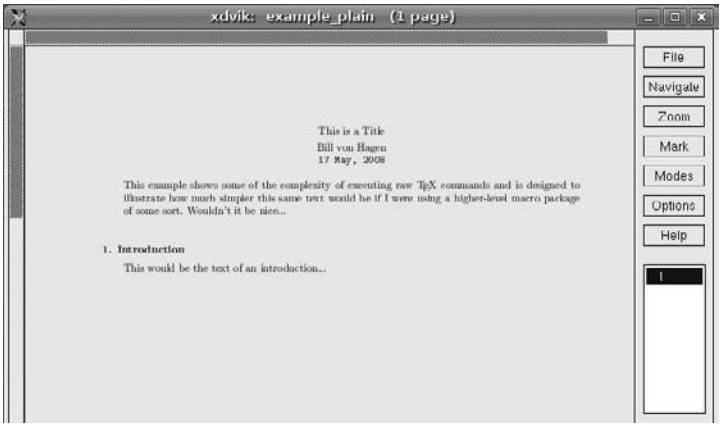
\bigskip

\beginsection 1. Introduction
\smallskip
This would be the text of an introduction...
\bye
```

Figure 16-1 shows a preview of the DVI file produced by processing this plain TeX document, shown in the `xdvi` DVI previewer that is installed as part of the `texlive` package.

FIGURE 16-1

Previewing output from the plain TeX example



Introducing LaTeX Markup

Let's face it—regardless of how good the output looks, writing documents using the low-level TeX commands shown in the previous section would be a tremendous pain, especially the extensive use of manual formatting and spacing commands. To simplify creating documents using TeX, a variety of macro packages have been developed, the best known of which is Leslie Lamport's LaTeX macro package. These macro packages support different types of documents, such as articles, manuals, letters, and so on, and provide easy-to-remember macros for the traditional components of those types of documents. For your convenience, the LaTeX macros package and associated applications are automatically installed when you install texlive on your Ubuntu system.

As an example of using the LaTeX article document type, the previous plain TeX code would look like the following in LaTeX:

```
\documentclass{article}
\begin{document}
```

```
\title{This is a Title}
\author{Bill von Hagen}
\date{\today}
\maketitle

\begin{abstract}
This example shows how much simpler this same document would become
when marked up using \LaTeX\ macros rather than raw \TeX\
commands, and is designed to illustrate how much simpler this same text
would be if I were using a higher-level macro package of some
sort. It sure is nice (and can also be understood by mere mortals).
\end{abstract}

\section{Introduction}
This would be the text of an introduction...
\end{document}
```

Running LaTeX produces slightly more verbose output than the previous example of running TeX. The output from processing the previous example, stored in the file `example_latex.tex`, is the following:

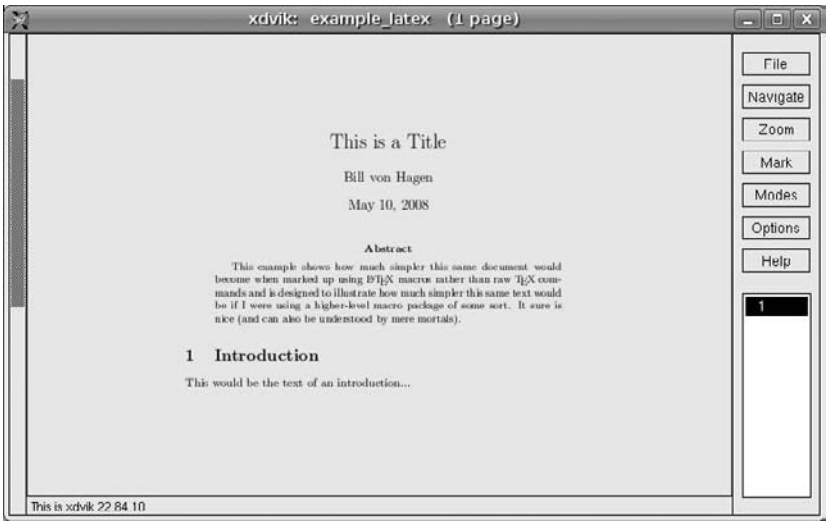
```
$ latex example_latex.tex
This is pdfTeXk, Version 3.141592-1.40.3 (Web2C 7.5.6)
 %&-line parsing enabled.
entering extended mode
(./example_latex.tex
LaTeX2e <2005/12/01>
Babel <v3.8h> and hyphenation patterns for english,
    usenglishmax, dumylang, nohyphenation, loaded.
(/usr/share/texmf-texlive/tex/latex/base/article.cls
Document Class: article 2005/09/16 v1.4f Standard LaTeX document class
(/usr/share/texmf-texlive/tex/latex/base/size10.clo))
(./example_latex.aux)
[1] (./example_latex.aux) )
Output written on example_latex.dvi (1 page, 1172 bytes).
Transcript written on example_latex.log.
```

Note that the command to format a LaTeX document is `latex` rather than the `tex` command used earlier. The `latex` command invokes a version of TeX that has been precompiled with all of the macros, font data, and other information that LaTeX defines. Figure 16-2 shows a preview of the DVI file produced by processing this LaTeX document, shown in the `xdvi` DVI previewer. As you can see from the sample LaTeX input, LaTeX provides macros for logical elements such as title, date, an abstract, sections (with automatic numbering), and so on. Figure 16-2 shows that the output is essentially identical, although with some differences in margins, headings, and so on.

Most people who create documents using TeX actually use LaTeX nowadays for the convenience of its macros, and also because it supports the automatic creation of things such as a table of contents, lists of figures, an index, and other document constructs that are both important and useful. Discussing the details of creating LaTeX documents is outside the scope of this book, but many other people have already done that. For suggested sources of detailed information about creating LaTeX documents and using LaTeX, see "More Information About TeX and LaTeX" later in this chapter.

FIGURE 16-2

Previewing output from the simple LaTeX example



Formatting and Previewing DVI Files

As mentioned previously, formatting a TeX or LaTeX document using the `tex` or `latex` commands produces a DVI file as output. You can instantly preview your DVI files using an application called xdvi, or you can use DVI postprocessors to generate output that is suitable for printing on the type of printer that you are using. Figures 16-1 and 16-2 show DVI files being previewed in the xdvi application. (The `texlive` package actually installs xdvik, which is a version of xdvi that supports hyperlinks and other modern conveniences.) The DVI postprocessors that produce output for various printers are as follows:

- `dvi2fax`: Produces output in G3 fax format.
- `dvihp`: Produces standard Hewlett-Packard PCL (Printer Control Language) output.
- `dvilj`: Produces output suitable for a Hewlett-Packard LaserJet.
- `dvilj2p`: Produces output suitable for a Hewlett-Packard IIp LaserJet.
- `dvilj4`: Produces output suitable for a Hewlett-Packard IV LaserJet.
- `dvilj6`: Produces output suitable for a Hewlett-Packard VI LaserJet.
- `dvips`: Produces output suitable for any PostScript printer.
- `dvipdf`: Produces PDF output that you can view using Adobe Acrobat Reader, the c PDF viewer, or any other PDF viewer.
- `dvired`: Produces 2-up output suitable for any PostScript printer. (This application requires `pstops`, which must be separately installed from the `psutils` package.)

All of these applications were installed on your system as part of the basic texlive installation described in "Installing TeX and LaTeX."

TIP **If you need to use TeX or LaTeX for the documents that you're working on, or are simply interested and want to experiment, some additional DVI-related utilities are available in the `dviutils` package, which you can install using package management software such as the Synaptic or Adept package managers. The `dviselect` utility enables you to print selected pages from or organize the pages in a DVI file, while the `dviconcat` utility enables you to combine DVI files and print them 2-up (two DVI pages per printed page) to save paper or simplify review.**

Using Graphical Tools to Work with TeX and LaTeX

As much of this book shows, it's a graphical world nowadays, and the Ubuntu and Kubuntu distributions provide all of the graphical applications that any computer user expects to find. Although command-line utilities such as TeX and LaTeX are easy to use, and the text-based format of TeX and LaTeX files makes them easy to exchange with users of many other computer systems and applications, it's hard to beat the convenience of a graphical application, especially when you're doing something such as word processing.

It's a good thing that the folks in the world of open source have stepped up to provide graphical solutions for creating and working with documents in LaTeX format. These are the Kile (the KDE Integrated LaTeX Environment), Lyx, and GNU TeXmacs applications, all of which are available in the Ubuntu repositories. Like texlive, these packages are not installed by default on your Ubuntu or Kubuntu system, but are easy enough to locate and install using the Synaptic Package Manager or Adept Manager. (See Chapter 23, "Adding, Removing, and Updating Software," for more detailed information on installing applications using these tools.) Installing Kile adds the Kile - LaTeX Frontend command to the Office menu. Installing Lyx adds the `Lyx` Document Processor command to the Office menu. Installing TeXmacs doesn't add anything to your menus, so you will have to execute TeXmacs by running the `texmacs` command from a command line.

CAUTION **Kile is a KDE-based application, so installing it on a standard GNOME-based Ubuntu desktop system will also install all of its KDE dependencies, but disk space is cheap and Kile is an excellent LaTeX editing and development environment.**

If you are working with LaTeX documents, it's hard to beat the convenience of Kile or the Lyx document processor. Kile is a KDE-based application that provides an amazingly complete editing environment for standard LaTeX documents, providing wizards for creating various document elements (the table wizard is especially impressive), easy searching for document elements (and where they end), and a very cool LaTeX menu from which you can select and insert most LaTeX constructs with a simple mouse-click. If you are new to LaTeX or simply want a sophisticated graphical environment for creating and working with LaTeX documents, I heartily recommend Kile, regardless of whether you're using Ubuntu or Kubuntu. Installing its KDE requirements on a standard GNOME-based Ubuntu system is well worth it, in my opinion. Figure 16-3 shows the same example LaTeX document in Kile.

As a KDE application, Kile is configured to use other KDE-based TeX and LaTeX applications. The most important of these is `kdvi`, the KDE DVI file previewer. Although you can configure Kile to use the standard xdvi DVI previewer (by choosing the Settings ➪ Configure Kile menu entry and modifying the command used to preview DVI files that is defined in Tools ➪ Build ➪ ViewDVI), you may simply want to install the `kdvi` package on your system if you are using Kile. Figure 16-4 shows our sample LaTeX document being previewed in `kdvi` after having been formatted using Kile's Build ➪ QuickBuild command.

FIGURE 16-3

Editing a LaTeX document in Kile

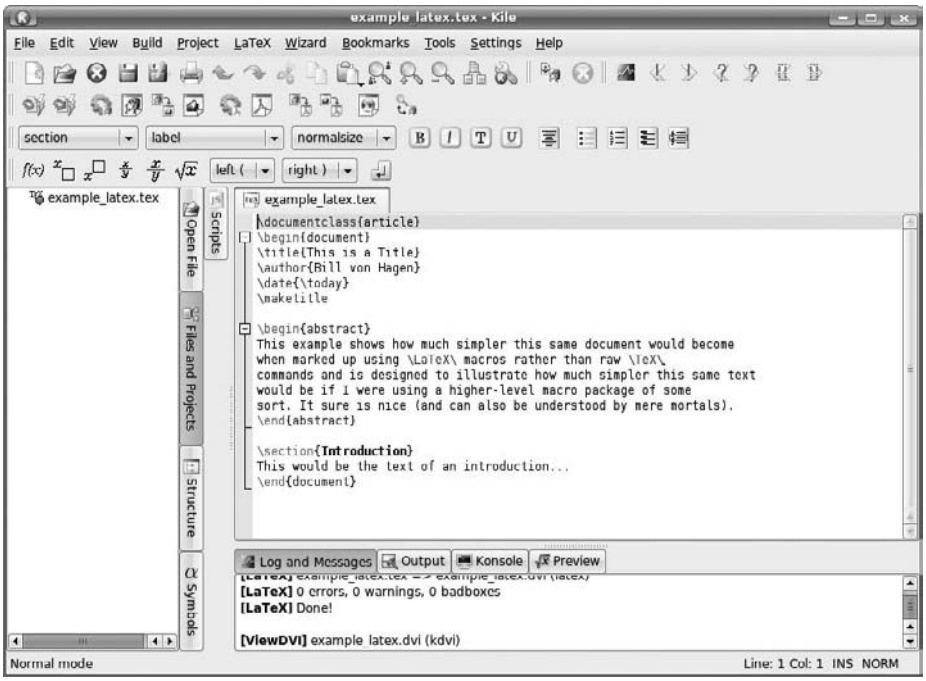
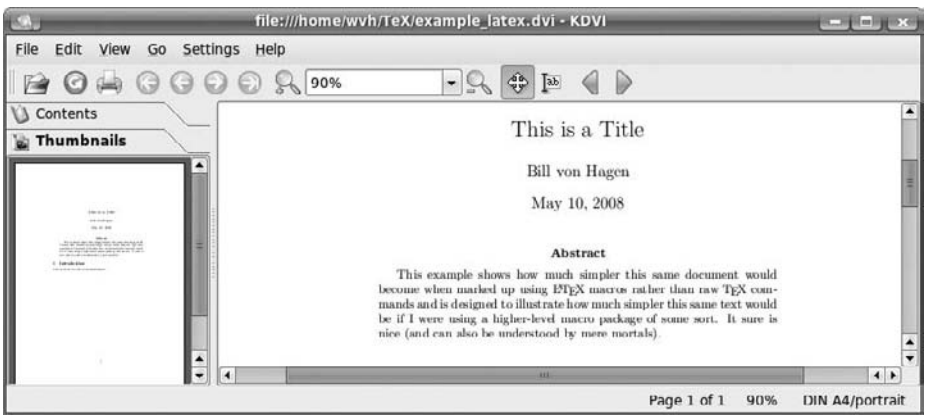


FIGURE 16-4

Previewing a DVI file in kdvi



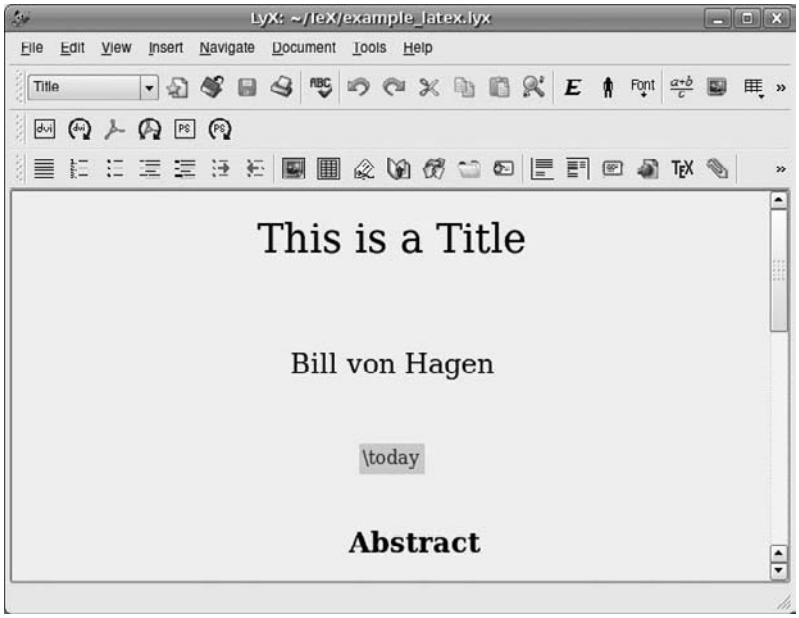
Lyx provides a convenient graphical interface for creating and working with LaTeX documents, including conveniences such as being able to select document elements and styles from drop-down menus and reformat selected characters, paragraphs, and so on, all without having to remember the low-level LaTeX commands that perform these functions in your text.

One of the coolest things about Lyx is its WYSIWYG (What You See Is What You Get) display and editing environment. Figure 16-5 shows the LaTeX document that is used as an example throughout this section in the Lyx document processor. You'll immediately notice that the onscreen representation of this document is almost identical to the appearance of the DVI file that was produced by formatting it using LaTeX. Maybe quasi-WYG would be a better term than WYSIWYG, but that term has never caught on.

To facilitate previewing and simplify editing, Lyx actually uses its own format for creating and storing documents, known as lyx format. To work with standard LaTeX documents, you must import them into Lyx using the File menu's Import command, and you can export them as standard LaTeX once you're ready to share them with your co-authors or reviewers. If you compare them with any LaTeX file that you originally imported, you'll note that Lyx adds a few more detailed LaTeX formatting commands to its LaTeX output, but the LaTeX files produced by Lyx are completely normal otherwise.

FIGURE 16-5

Editing a LaTeX document in Lyx



Lyx is an excellent tool, and well worth experimenting with if you are working with LaTeX documents but want the convenience of a graphical application or simply can't remember all those pesky LaTeX commands. For more information about Lyx, see its home page at `www.lyx.org`.

Another alternative for working with LaTeX documents in a graphical form is the TeXmacs package, which is a graphical, LaTeX-aware package that is implemented on top of GNU emacs. This enables you to use familiar Emacs commands while working with your LaTeX documents in a graphical fashion. Figure 16-6 shows the same LaTeX document used as an example throughout this chapter in the TeXmacs document processor. You'll immediately notice that the onscreen representation of this document is almost identical to the appearance of the DVI file that was produced by formatting it using LaTeX.

After installing TeXmacs, you must execute it from a Linux command line by typing the `texmacs` command, or selecting the Accessories ➪ GNU TeXmacs Editor menu command. If you execute TeXmacs from the command line, you'll see that it does a lot of initial setup and configuration, which you can ignore—and don't worry, it will be quiet the next time you run it.

TeXmacs is convenient to use and provides several LaTeX-specific menus and icons that make it easy to apply LaTeX styles, create document elements, and even select and reformat portions of your documents. Unlike Lyx, you can use TeXmacs with TeX documents as well as LaTeX documents, and it is not necessary to import your LaTeX documents—TeXmacs will open your existing LaTeX files, but they may require a bit of tweaking to get them to display correctly. TeXmacs will also offer to save your LaTeX files in its own format, using the `.tm` file extension. For more information, see the TeXmacs home page at `www.texmacs.org` or consult the extensive online help that it provides.

FIGURE 16-6

Editing a LaTeX document in TeXmacs



TIP **If you want to work on your LaTeX files in a graphical environment, I think that Kile is the best of the applications discussed in this section, without question. While Kile isn't WYSIWYG (or even quasi-WYG), its power and ease-of-use more than make up for that, especially if you are new to LaTeX. If you're using a standard GNOME-based Ubuntu desktop system, you may flinch at having to install tons of KDE dependencies in order to use Kile, but my job is to suggest the best solutions to different problems, not necessarily to save you disk space. Given the size of today's hard drives, what's a few gigabytes between friends, anyway?**

More Information About TeX and LaTeX

Because TeX and LaTeX have been used for many years on a wide variety of systems, a tremendous amount of information about using them is available on the Web. A simple Internet search will turn up more hits than you could possibly want. As a timesaver, some of my favorite sites for information about TeX and LaTeX are as follows:

- Getting to grips with LaTeX (`www.andy-roberts.net/misc/latex/index.html`) is a great series of tutorials on using LaTeX from Andrew Roberts.
- Introduction to LaTeX (`www.math.uiuc.edu/~hildebr/tex/course`) is a short course prepared at the University of Illinois Department of Mathematics that provides a nice, quick overview of TeX and LaTeX.
- Text Processing with LaTeX (`www.eng.cam.ac.uk/help/tpl/textprocessing`) is a great site at Cambridge University that provides many online tutorials and examples of creating a variety of documents using LaTeX.
- The teTeX HOWTO (`www.tldp.org/HOWTO/TeTeX-HOWTO.html`) discusses the previous version of TeX/LaTeX that was used on Ubuntu and Kubuntu systems. This HOWTO was written a few years ago, but one of the nice things about stable software like TeX and LaTeX is that things don't change much. This document provides a good deal of general information about creating TeX and LaTeX documents, but you should ignore any specific information about using the teTeX TeX implementation.

Several books dedicated to TeX and LaTeX are available and a quick search at Amazon.com will provide an assortment of options. However, no discussion of TeX and LaTeX would be complete without identifying the two seminal texts for each of these:

- *TeXbook* by Donald E. Knuth (Addison-Wesley Professional, 1984, ISBN: 0201134489)
- *LaTeX: A Document Preparation System*, 2nd Edition by Leslie Lamport (Addison-Wesley Professional, 1994, ISBN: 0201529831)

These books are by the authors of the respective systems, you can consider them definitive. Some changes have certainly been made to LaTeX since its creation by Lamport long ago, but these are still the bibles for both of these document preparation systems.

Word Processing with OpenOffice.org Writer

OpenOffice.org (`www.openoffice.org`) is a suite of office applications including a word processor (Writer), a spreadsheet (Calc), a drawing program (Draw), software for creating presentations (Impress), and a database program called Base, which keeps tagging along after the rest of OpenOffice. This section focuses on the word processing application—other chapters in this book discuss the OpenOffice.org spreadsheet, graphics, and presentation software. But first, a word from our sponsor.

OpenOffice.org is one of the most powerful and popular open source projects on the Web, and the people who are largely responsible for its existence are the folks at Sun Microsystems. Scott McNealy, president of Sun Microsystems, has a long-standing rivalry with Microsoft, so it wasn't all that surprising when in 1999 Sun purchased StarDivision, the German vendor of a suite of office software called StarOffice.

The purchase of StarOffice was originally seen as both a way to leverage Sun's Java offerings (because StarOffice and OpenOffice.org are both heavily dependent on Java), and to provide Sun users with an office suite that they could run natively on their workstations.

Sun then released the then-current source code for StarOffice into the open source community under the terms of Sun Industry Standards Source License (SISSL). Beginning with release 2.0 of the OpenOffice.org suite, all of the source code is now available under the GNU Lesser General Public License (LGPL). Sun still markets StarOffice, which is both supported and offers some capabilities beyond those provided by OpenOffice.org—but that's moot for us, because StarOffice doesn't run on Ubuntu Linux.

Aside from its power and the absence of cost, a truly compelling reason for using OpenOffice.org Writer is its use of the new OpenDocument Format (ODF) standard for storing and exchanging documents (`www.oasis-open.org/committees/tc_home.php?wg_abbrev=office`). Historically, all word processors used their own proprietary binary formats for storing documents, which made it a hassle when switching from one word processor to another or when simply trying to exchange documents with friends or coworkers who used a different word processor. The OpenDocument standard is designed to provide a published standard for storing documents in an XML format that can instantly be used in any other word processor that supports this format. This may seem like a minor thing, but get back to me after you try to open the WordStar files for your family history, short stories, or other documents that you may have written long ago but still want to be able to use.



If you are interested in document standards or read nerd news occasionally, you may be aware of the arm-wrestling contest between the ODF and Microsoft's OOXML (Open Office XML) standards regarding which of these gets to be the official open document standard. OOXML "won" this battle, but in my opinion you would be foolish to use OOXML. A standard that is over 5,000 pages long (and which not even the Microsoft tools conform to at the time of this writing) isn't much of a standard as far as I'm concerned. Use ODF and reject this latest offering from Microsoft. It's one thing that their tools and operating system are bloatware—it's another to make bloat the standard.

To avoid repeating the phrase "OpenOffice.org Writer" more times than necessary throughout the rest of this section, I'll typically refer to OpenOffice.org Writer simply as Writer (its name to its friends) or occasionally as `oowriter` (the name of the application that you must type to execute OpenOffice.org Writer from a command-line interface).



At the time of this writing, version 3.0 of the OpenOffice.org suite is available, but has not yet been integrated into Ubuntu or Kubuntu. OpenOffice.org v3 adds many new features, the most important of which may be its support for the latest Microsoft "standard" document and other file formats. I'm sure that they'll be integrated soon, but if they haven't been integrated by the time that you read this, you can install them yourself. To install these, use the sudo command and your favorite text editor to add the following line as the last line of your `/etc/apt/sources.list` file:

```
deb http://ppa.launchpad.net/openoffice-pkgs/ubuntu \
intrepid main
```

(You should enter this on a single line, with no backslash). After saving the file and exiting the editor, update the list of available packages and sources in whatever package management software you are using. You will then be able to install the OpenOffice.org 3.0 packages as part of your next update.

File Extensions for OpenOffice.org Documents

With the adoption of the OpenDocument format, OpenOffice.org Writer changed the names of its default file extensions. The following list shows the file extensions used by different OpenOffice.org versions:

- `odm`: OpenDocument Master Document
- `odt`: OpenDocument Document, OpenOffice.org 2.0 and greater
- `ott`: OpenDocument Document Template, OpenOffice.org 2.0 and greater
- `sxw`: Writer Document before OpenOffice.org 2.0
- `stw`: Writer Document Template before OpenOffice.org 2.0
- `sxg`: Writer Master Document before OpenOffice.org 2.0

The most common documents that you will see and create are those with the ODT and OTT file extensions. However, you may encounter the older SXW format when exchanging documents with long-time OpenOffice.org users. SXW files are very interesting, because they are actually zipped archives containing XML document files and associated metadata. You must use the Document Converter Wizard to import old-format SXW OpenOffice.org files into the new OpenOffice.org OpenDocument format, as explained in the section "Importing Documents from Other Word Processors" later in this chapter.

Installing Files for Writer

Writer is installed by default if you install Ubuntu or Kubuntu from a Live CD, from an Alternate Install CD using the Text Mode installer, or from an Alternate Install CD using the OEM Mode installer. If you want to add Writer or the entire OpenOffice.org suite to a server system that has a graphical user interface, you can install them using `apt-get`, `aptitude`, or `Synaptic` (or the `Adept Manager` on Kubuntu systems), as explained in Chapter 23, "Adding, Removing, and Updating Software."

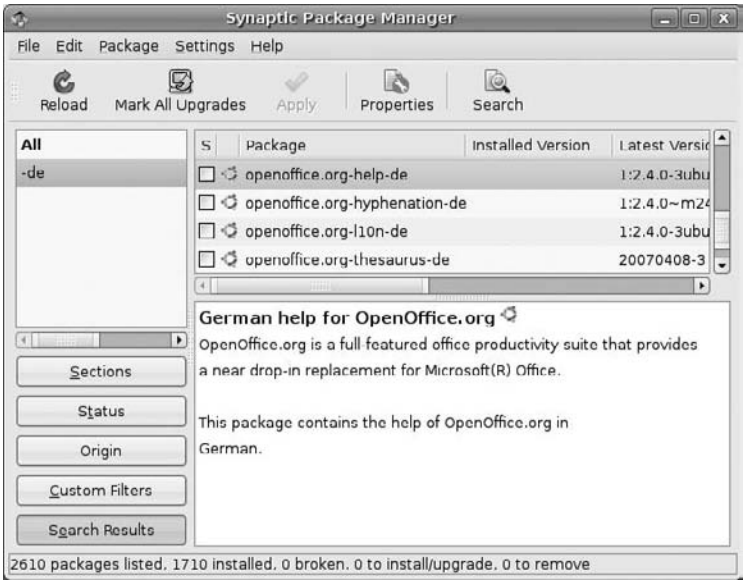
One thing that differentiates Ubuntu and Kubuntu from other Linux distributions is their extreme dedication to the user, especially in terms of internationalization. If you are working in only a single language, the correct localization packages for your physical location are installed with your default Ubuntu/Kubuntu installation. However, if you are installing Writer for use when creating documents in multiple languages, you may also want to search for and install the appropriate hyphenation, localization, and thesaurus packages for whatever other locales you are writing for. You can easily find these by looking for a string that represents the country code that you are looking for. For example, Figure 16-7 shows the specific German language support packages for OpenOffice.org that have been identified by searching for the string "-de" in the Synaptic Package Manager.

TIP **If you've already installed the localization and internationalization packages for any component of the OpenOffice.org suite, there is no need to reinstall them. Installing them once installs them as a shared resource that can be used by every component of the suite.**

TIP **You can install complete language support for all applications (without translating GNOME or KDE themselves) by installing the appropriate language support metapackage for your target locale(s)—these packages have names of the form language-support-country-code (for example, language-support-de), and include the dictionaries for the specified locale. To fully translate the GNOME or KDE interface and all internationalized applications, select the appropriate language pack. The names of the language packs are of the form language-pack-country-code. Once you've installed multiple dictionaries, and so on, you can set the default language for your documents in Open Office applications using the Tools ➪ Options ➪ Language Settings ➪ Language drop-down list. You can also set the default language for various portions of your documents by using the commands on the Tools ➪ Language menu.**

FIGURE 16-7

Sample additional language packages for the OpenOffice.org suite



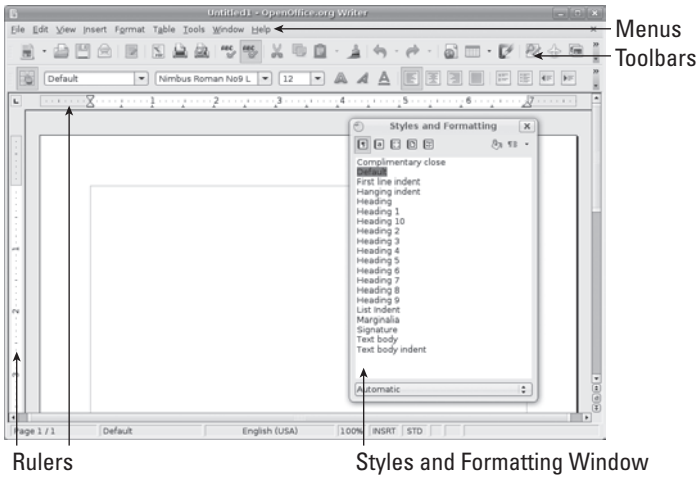
See Chapter 23, "Adding, Removing, and Updating Software," for more detailed information on installing applications using the Synaptic Package Manager, the Adept Manager on Kubuntu systems, and related command-line utilities.

Taking a Quick Tour of Writer

You can start Writer from the command line by executing the `oowriter` command. More conveniently, a menu item for Writer is installed in the Applications ➪ Office menu called OpenOffice.org Word Processor. To start Writer from the menus, select this menu item. The application starts, displaying an empty document, as shown in Figure 16-8.

FIGURE 16-8

Starting Writer

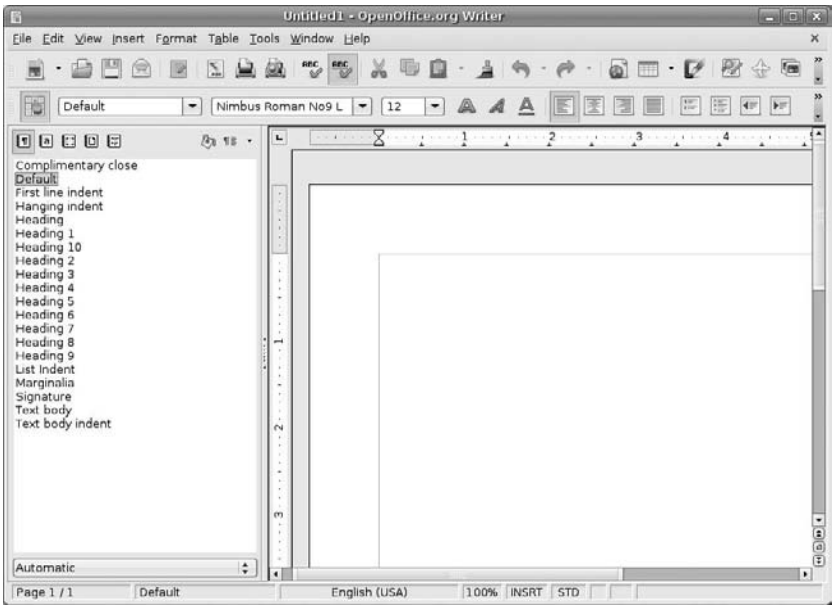


Writer provides excellent online help, as do all of the tools in the OpenOffice.org suite, so I won't bore you by walking you through each menu and entry that it contains. However, as a quick overview of the Writer interface, the callouts in Figure 16-8 highlight the following areas:

- **Menus:** The drop-down menus in Writer are organized much like the menus in Microsoft Word to provide a familiar user experience. Menu items with an arrow at their right edge lead to other, related submenus.
- **Toolbars:** By default, Writer displays the Standard and Formatting toolbars, which provide one-click access to specific commands when you click on the appropriate icon. You can modify the list of icons displayed on a toolbar by clicking on the control for that toolbar, located at the far right. You can display other toolbars or deactivate the default toolbar by selecting toolbars from the View ➪ Toolbars submenu.
- **Rulers:** These rulers display the horizontal and vertical location of the text insertion point in the current document. The grayed-out portions of each ruler reflect the page margins used in the current document. You can hide these rules by toggling the View ➪ Ruler menu command.
- **Styles and Formatting window:** Although it's not displayed by default, you can display this window at any time by pressing the F11 key on your keyboard. (Press F11 again to close this window.) The Styles and Formatting window provides easy access to various sets of styles that are available for the current document, and is especially convenient if you choose not to display the formatting toolbar. The drop-down menu enables you to display specific sets of styles in this window, such as Character Styles, List Styles, and so on. You can dock this window by holding down the Ctrl key and double-clicking on the gray portion of the window heading (beside the icons), making it a part of the default Writer interface rather than a floating window. Figure 16-9 shows Writer with the Styles and Formatting window docked. You can undock a docked window by repeating this process in the docked area.

FIGURE 16-9

Docking the Styles and Formatting window



- **Context-sensitive menu:** Right-clicking on the background of any document that you are working on in Writer displays the context-sensitive menu shown in Figure 16-10. This menu provides quick access to several formatting and style-related submenus.

Writer provides extensive online help, available by selecting OpenOffice.org Help from the Help menu, or by pressing the F1 key on your keyboard at any time. Figure 16-11 shows the main page for Writer's help, which provides both a searchable index and a list of popular, general topics from its initial display window (at right).

Personalizing Writer

The first thing that you'll want to do when getting started with Writer is to enter some information about yourself into its User data fields. This is primarily useful because Writer will use this information in various places when filling out fields in documents that you create using the OpenOffice.org Wizards (as explained in the next section).

Select the Tools ➪ Options menu item to display the Options dialog's User Data form, as shown in Figure 16-12. Fill out the fields in this form and Click OK to return to Writer and your document.

FIGURE 16-10

Writer's context-sensitive menu

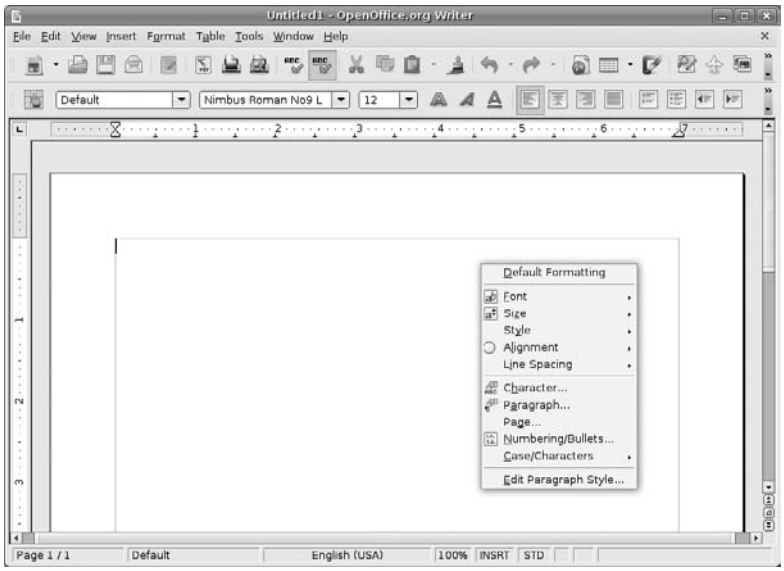


FIGURE 16-11

Writer's Online Help

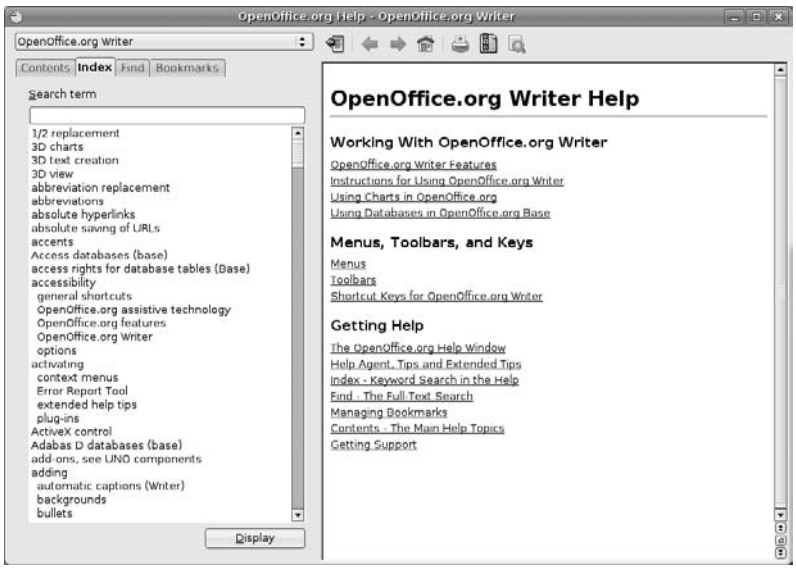
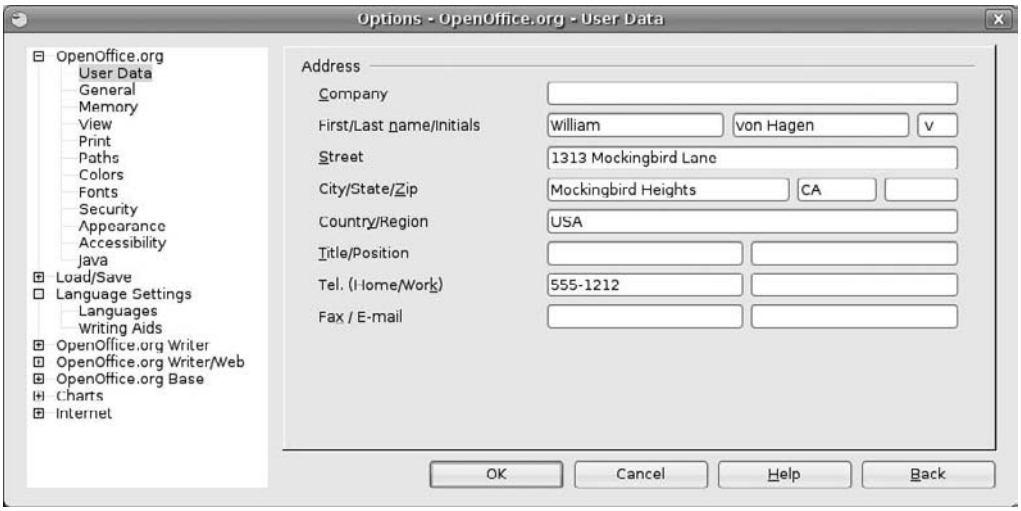


FIGURE 16-12

Writer's User Data form



Using Wizards to Create Documents

The easiest way to create a document in Writer is to use one of the document wizards that it provides for many popular types of documents. This section provides an example of creating a letter in Writer using the Letter wizard, but the process of using any wizard is quite similar regardless of the type of document you're creating—and the goal of these wizards is to be easy to use in the first place.

To start the appropriate wizard for creating a new document, do the following:

1. Select File ➪ Wizards, and select the appropriate wizard for the type of document that you are creating. Writer provides wizards for creating letters, faxes, agenda, presentations, and Web pages. For this example, I'll select the Letter wizard. Writer creates a new letter document using its default, and displays the first Letter Wizard dialog, as shown in Figure 16-13.
2. Select the radio button for the type of letter that you want to create: business, formal personal, or personal. Each type of letter provides some default page styles and decorations, which you can select from the Page design drop-down list. The window behind the Letter Wizard shows a preview of a letter with your current selections, so that you know what you're getting into. If you look closely, you'll see that the Letter Wizard has already filled out some fields in the letter from the User Data form that was discussed in the previous section. In this example, I'll select Business Letter. Once you've made your selections, click Next to proceed. The second Letter Wizard dialog displays, as shown in Figure 16-14.

TIP **If you are creating a business letter and you are going to be printing it on your corporate letterhead, you can select the Use letterhead paper with pre-printed elements checkbox (shown in Figure 16-13). This removes any existing graphics from the letter preview, and activates the Letterhead layout step in the wizard, so that you can specify the dimensions of your letterhead stationery and thus avoid overprinting portions of it with your letter.**

FIGURE 16-13

The first Letter Wizard dialog

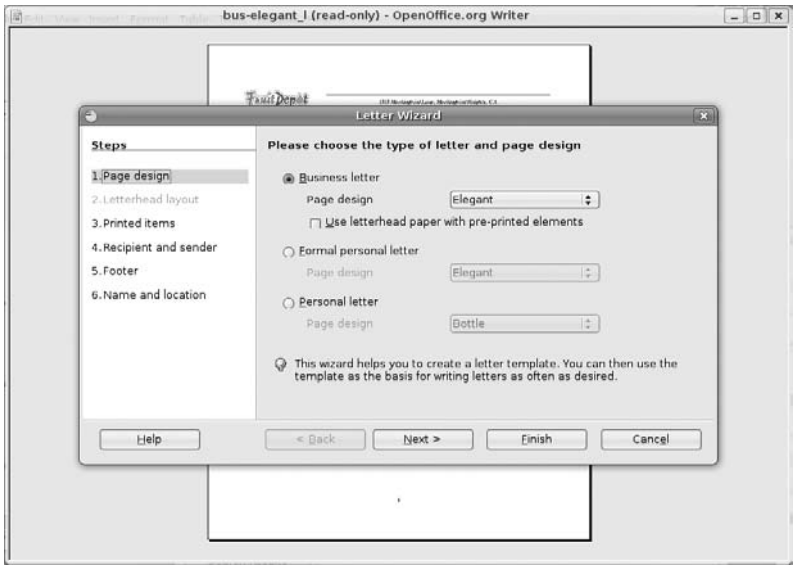
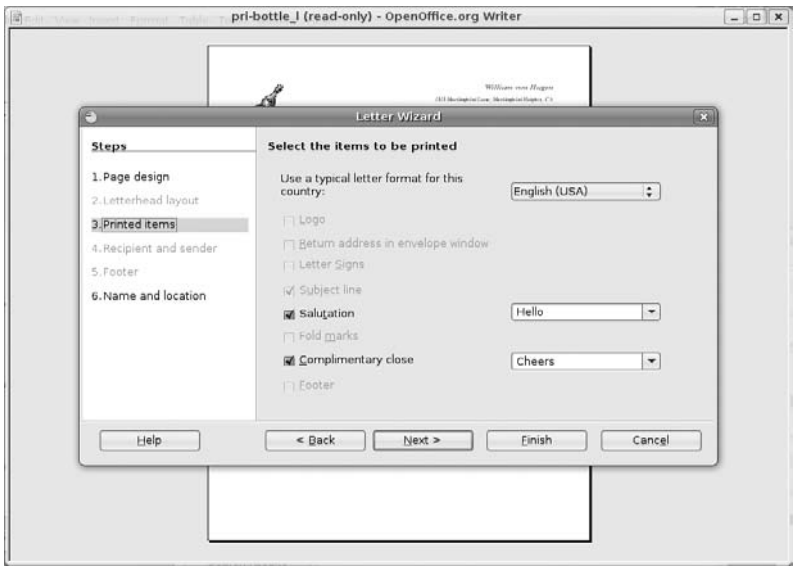


FIGURE 16-14

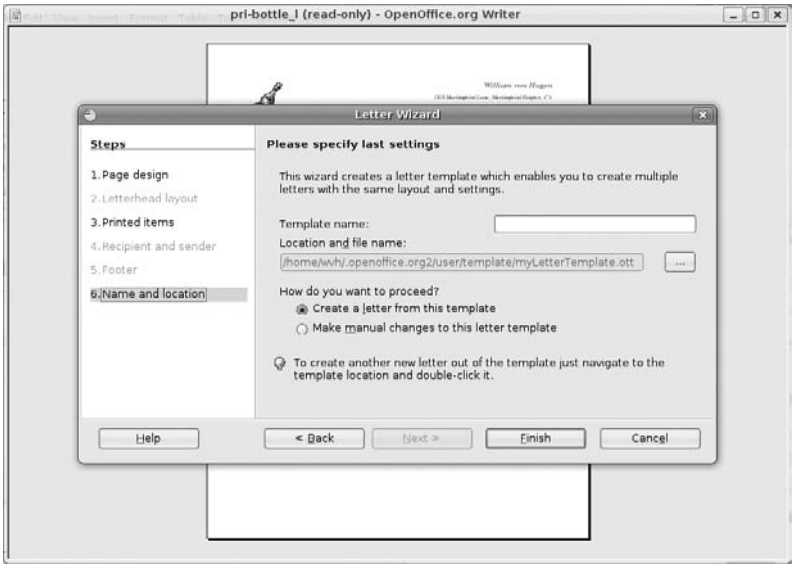
The second Letter Wizard dialog



3. Select the Salutation or closing message for your letter. You can change these by selecting existing values from the drop-down menus or personalize them by selecting and replacing the default values shown for these. As you update these, the preview of your letter changes to reflect the new values. When you're done, click Next to proceed to the next screen. The Letter Wizard's third (and final) dialog displays, as shown in Figure 16-15.

FIGURE 16-15

The final Letter Wizard dialog



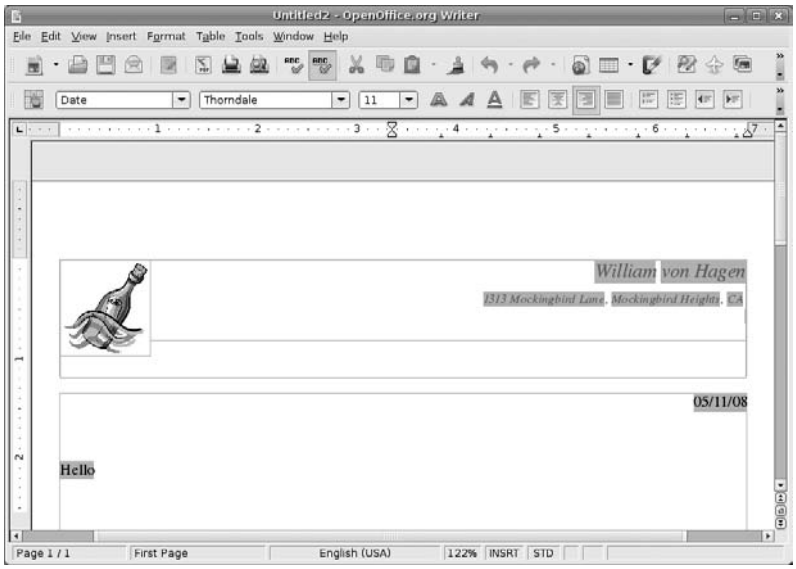
4. If you've made any changes to the default letter settings and are going to be reusing them, the final Letter Wizard dialog enables you to update the wizard's default settings or save your modifications as a new template. To update your default settings, enter a name to associate with your new settings. If you want to create a new template from these settings and continue writing your letter, click the Browse button and enter a new filename for the new template. This file will still be used the next time you use the Letter Wizard, but you can easily create documents from different templates as described later in this chapter in the section "Creating and Using Templates."

Once you're done, click Finish to terminate the Letter Wizard and proceed with actually writing your letter. Your new letter displays in the main Writer window, as shown in Figure 16-16.

TIP **If you'd rather tweak the template further before actually using it, select the radio button beside "Make manual changes to this letter template" (shown in Figure 16-15) and click Finish. This displays the template, at which point you can add or modify styles and page formatting, as described in the next section.**

FIGURE 16-16

Editing your new letter



As you can see, Writer's wizards make it easy for you to create documents of certain predefined types. Writer's wizards leverage some of the predefined document templates that are provided with a default installation. If you want to create your own documents based on a template that is not supported by a wizard, Writer still makes it easy for you to define and use your own, specialized types of documents, as described later in this chapter in the section "Creating and Using Templates."

Modifying Document Styles and Layout

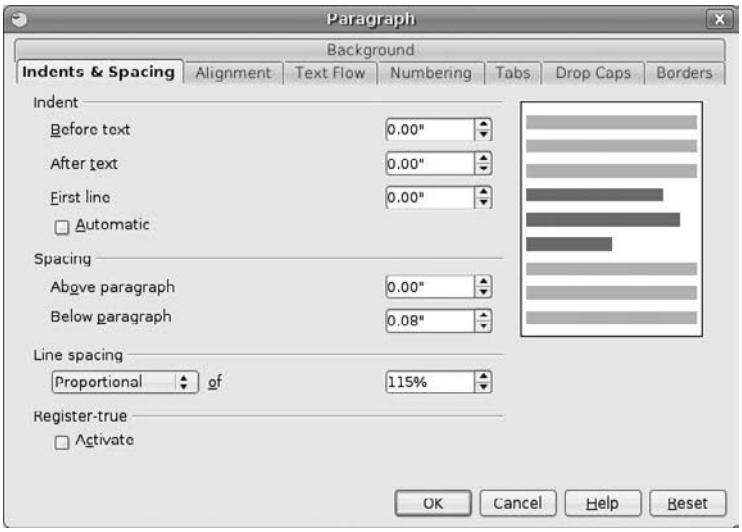
A style is the word processing term (borrowed from more advanced page layout programs like InDesign and PageMaker) for a set of formatting and display characteristics that are associated with a specific portion of your text, usually a paragraph. Once you're actually editing a document, you can apply any of the existing styles that have been defined for that document by selecting the name of a style from the drop-down list in the formatting window, or by selecting it from the Styles and Formatting window. To display the Styles and Formatting window, press F11. (Press F11 again to close the Styles and Formatting window.)

All of Writer's default document templates provide a default set of styles that enable you to format the different portions of your document. To apply a style to an existing paragraph, simply click anywhere inside it using the mouse and select the style that you want to apply from the drop-down list or Styles and Formatting window. I'm using the term "paragraph" here because, by default, a paragraph is any element in your document that ends with a hard return, even if it's empty.

To modify a single instance of a paragraph formatted with an existing style, position the cursor in that paragraph, right-click, and select Paragraph from the pop-up menu. This displays the dialog shown in Figure 16-17, which enables you to modify the style settings used in a single element in your document without changing them for other elements in your document that use the same style.

FIGURE 16-17

Modifying the style for a single paragraph



Although all OpenOffice.org documents come with a good set of default styles, everyone's tastes differ. To modify an existing style to reflect your personal tastes or corporate standards, right-click on any portion of your document that is currently formatted with that style, and select Edit Paragraph Style from the pop-up menu. A dialog like the one shown in Figure 16-18 displays.

Changing any of the settings in this dialog will change the style itself, which means that any other elements of your document that use this style will also be updated with your modified settings. You'll notice that the dialog shown in Figure 16-18 has many more options and tabs than the simple Paragraph formatting dialog shown in Figure 16-17. This is because the dialog in Figure 16-18 enables you to modify any aspect of the style, not just the formatting of a selected element.

Changing the default styles is fine, but may be the wrong thing to do if you want to modify only a few elements in your document that use a certain style. When the default styles aren't sufficient and it's a waste of time to tweak multiple paragraphs, the right thing to do is to create a new style, based on an existing one, which has slightly different characteristics.

To create a new style based on an existing style, first make sure that the Styles and Formatting window is shown on your screen (press F11 if it is not). Next, position the cursor in a paragraph formatted with the existing style, and click the New Style from Selection icon at the far right of the Styles and Formatting window toolbar. Select New Style from Selection from the drop-down menu to display the dialog shown in Figure 16-19.

Enter the name that you want to associate with your new style, and click OK. Figure 16-19 shows me creating a style named Quote, which I use for indented quotations from other sources. Whatever value you enter, the name of your new style displays in the Styles and Formatting window, and the new style is associated with the currently selected element. You can then repeat the process for modifying an existing style, as described earlier in this section, to add the custom settings that you want to associate with your new style.

FIGURE 16-18

Modifying the values for an existing style

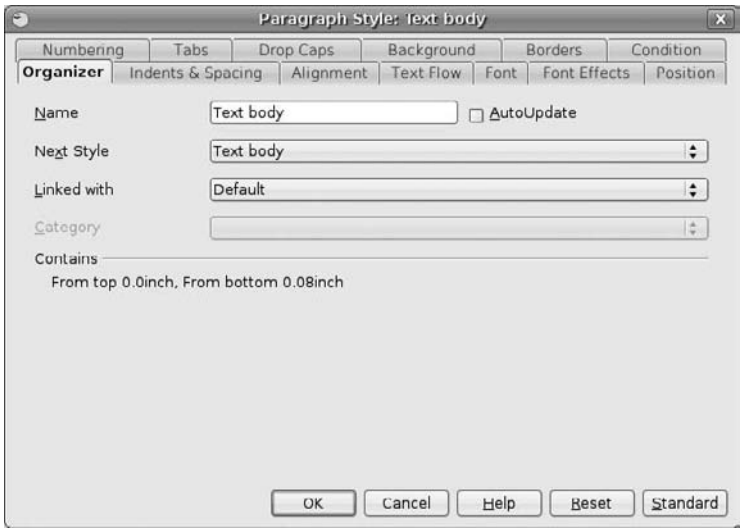


FIGURE 16-19

Creating a new style from a selection



TIP **You should, of course, minimize the number of new styles that you create whenever possible. For example, you may be working with corporate document standards that mandate specific formatting and styles. In this case, adding new styles may be, at worst, a violation of corporate policy and, at best, an irritation to anyone who subsequently has to work on a document that you've written or modified. Always make sure that it's absolutely necessary before creating a new style.**

Creating and Using Templates

As mentioned earlier in this chapter, Writer comes with a good selection of default document types, known as *templates*, and provides wizards to help you use them or even create new templates based on

those document types. The previous section described how to create your own styles to further fine-tune the settings that you use in your documents. This is very convenient, but only modifies the styles associated with a specific document—the template from which that document was created doesn't pick up the new styles, which makes it hard to share your stylistic insights with others. It also makes it hard for you to reuse them yourself, unless you subsequently create every new document by copying an existing document that has the new styles, and then replace its contents. At best, this is a pain in the anatomy, and at worst, it is almost completely inefficient. The right solution is to be able to easily create and manage your own document templates. You can then reuse your styles by starting new documents from your updated template, and even share your Writer wizardry with other users simply by giving them copies of your templates.

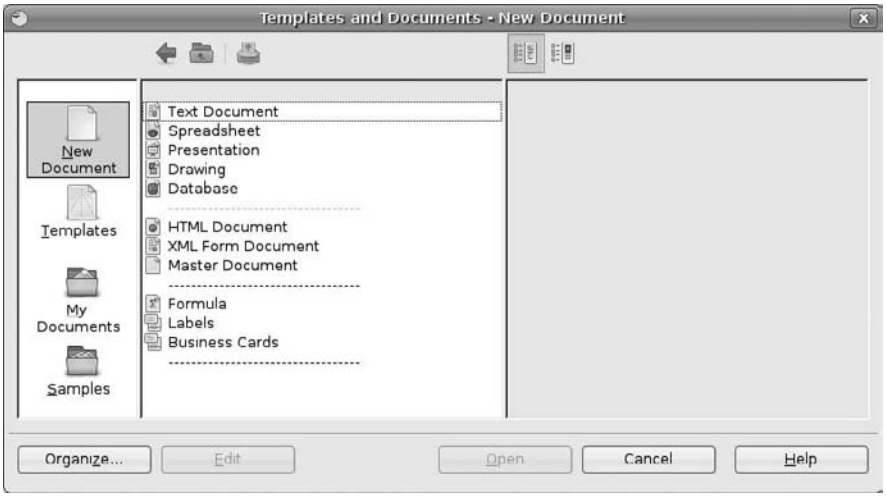
NOTE **The process of creating a new template is initially the same as that for manually creating a new document (although without using a wizard). The primary differences are that you don't need to add much content (only representative elements for designing new styles) and that you save the new document as a template rather than as a document. Read on!**

To create a new template, do the following:

1. Select File ➪ New ➪ Templates and Documents from any Writer window to display the dialog shown in Figure 16-20.

FIGURE 16-20

Creating a new template

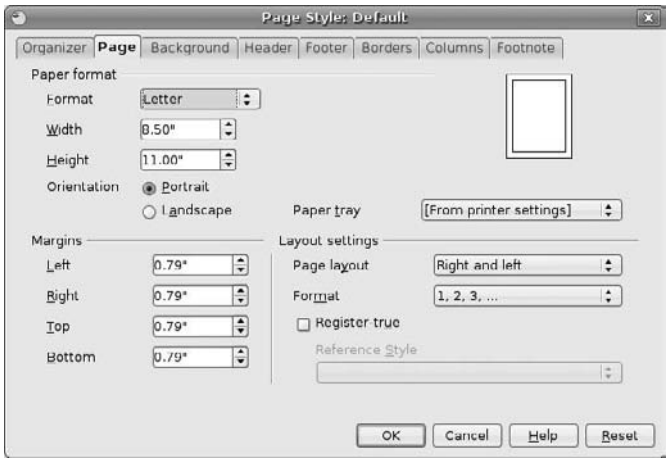


2. Select the New Document icon at the left, and select Text Document as the type of document that you're creating. Click Open to begin modifying the new document, which displays as an empty document in a new Writer window. In this window, create any stylistic elements and layout settings that you want to provide through your new template. (The previous section explained how to create new styles from existing ones in existing documents.)

3. Make sure that your page settings are correct for the new style. Beyond creating new styles, one common change that you will want to make in your new template is to change the default page layout by selecting Format ➪ Page, selecting the Page tab, and modifying the page size, margins, and other settings in the dialog shown in Figure 16-21.

FIGURE 16-21

Modifying page type and layout settings



4. Once you've finished modifying your new template, you can save it as a template by selecting File ➪ Templates ➪ Save. A dialog like the one shown in Figure 16-22 displays.

FIGURE 16-22

Saving your new template



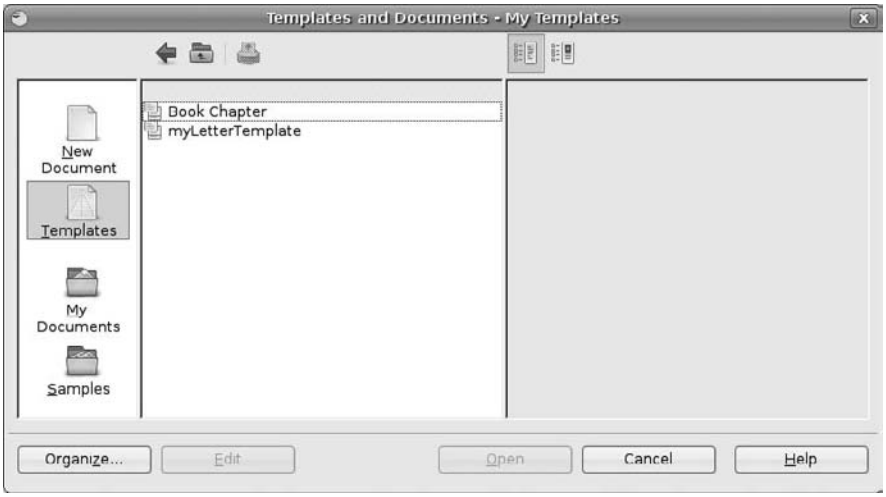
5. Enter the name for your new template and click OK. Figure 16-22 shows the creation of a new style called Book Chapter.

Congratulations, you've created a document template! Once you've created a template, you can easily create new documents based on that template by doing the following:

1. Select File ➪ New ➪ Templates and Documents to display the dialog shown in Figure 16-20.
2. To create a document from a specific template, click the Templates icon at the left and select the location where you saved the template (usually "My Templates"). You will see your templates listed in the window to the right of the icons. Figure 16-23 shows this dialog, which includes the template that I just created in the previous step.

FIGURE 16-23

Creating a new document from a template



3. Select the template that you want to use and either click Open or double-click on its name to create a new document from this template.

Once you are finished editing your document, save it by using the standard File ➪ Save As command.

Creating and using templates is a great way to standardize the documents that you or your organization is creating.

TIP **Modifying or otherwise updating a template does not mean that any documents that have been created from it are automatically updated. After updating a template, opening any document that was created from it will prompt you as to whether or not you want to update the document to reflect the updated template. That's usually a good idea, but it's nice that Writer is kind enough to ask you about it.**

Importing Documents from Other Word Processors

Writer can open documents produced by many other word processors, especially those created by Microsoft Word, Corel WordPerfect, or documents stored in the intermediate Rich Text Format (RTF)

document type. This is fine if you have only one or two documents that you want to update, but if you are moving all of your work to Ubuntu, you may find that you have hundreds of documents that you want to update to Writer format. It's a good thing that Writer provides a wizard for this that will not only import all of your documents, but also will import and convert any existing or attached style sheets that it detects.

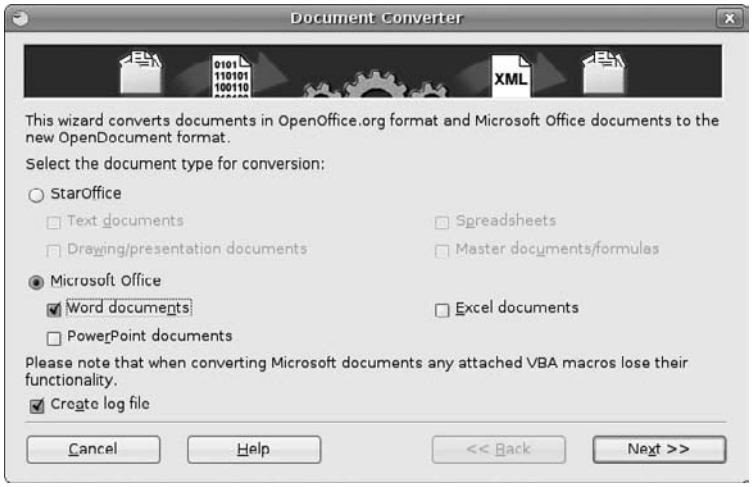
TIP **At the time of this writing, not all of the Java extensions used by the OpenOffice.org applications played well with various compositing effects provided by Compiz. If you have a problem seeing any of the dialogs from Writer or any other OpenOffice.org applications on a standard GNOME-based Ubuntu desktop system, you may have to use the Visual Effects tab on the Preferences ➪ Appearance dialog to temporarily change your effects preferences to None. These interaction problems will certainly be fixed soon, but at least there's a workaround until then.**

To batch-import your existing Microsoft Word (or older OpenOffice.org documents), do the following:

1. Select File ➪ Wizards ➪ Document Converter to display the dialog shown in Figure 16-24.

FIGURE 16-24

Specifying the input format when converting documents



2. To import Microsoft Word files and style sheets, select the Microsoft Office radio button and select the Word Documents checkbox. To import existing SXW documents from StarOffice or a version of OpenOffice.org prior to 2.0, leave the StarOffice radio button selected. Click Next to proceed. The dialog shown in Figure 16-25 displays.
3. Specify the directories in which you want the wizard to search for the specified document types, and the locations to which converted documents should be saved. You may want to update the location to which converted templates should be saved—I typically use my home directory, and subsequently move the files around manually. Once you've checked the directory locations, click Next to proceed. A status dialog displays, summarizing the conversions that are about to take place. Click Convert to begin the conversion process. The dialog shown in Figure 16-26 displays as the conversion process progresses.

FIGURE 16-25

Specifying document and template locations when importing

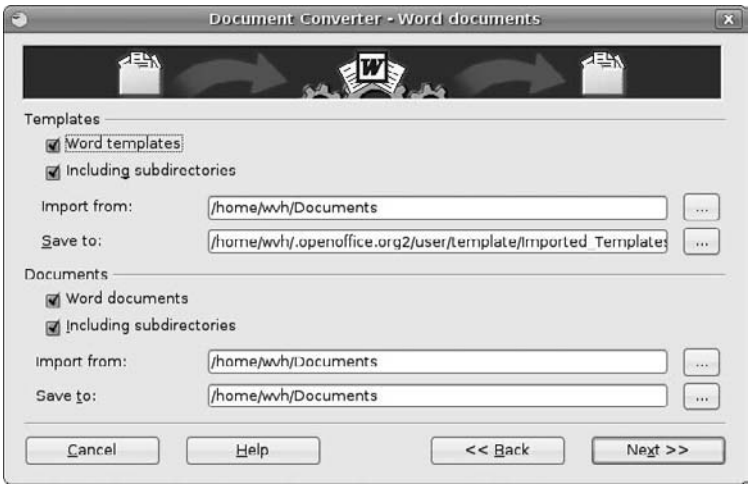


FIGURE 16-26

Monitoring the document conversion process



4. During the conversion process, the conversion wizard may display a character code confirmation dialog, asking you to select the character set, font, and language you want. This generally means that the wizard encountered unknown characters in one of the documents. When the conversion

process completes, the Close and Show log file buttons are displayed on the dialog shown in Figure 16-26.

Congratulations—the conversion process is complete! Your converted documents and templates are ready for use in Writer. Before clicking Close, it's a good idea to click the Show log file button. The log file for the conversion process lists successful conversions, and also identifies any documents that should be fine-tuned or corrected before they can be used with Writer. This is often the case with Word DOT files, which may contain macros that are not completely compatible with anything other than Microsoft's own Visual Basic for Applications (VBA).

More Information About Writer

As mentioned earlier in this chapter, all OpenOffice.org applications provide complete, well-written online help. However, like most online help, this information is reference information, not HOWTO information.

If you are serious about using OpenOffice.org Writer for all of your documentation needs, the best site that I have seen on the Web for HOWTO information about using Writer is `www.tutorialsforopenoffice.org`. This site features great tutorials in PDF and OpenOffice.org formats, as well as links to a great selection of other sites with useful, hands-on information about using Writer. Check it out—it's well worth your time! After you become an OpenOffice.org wizard, you may want to create a tutorial of your own and contribute it to the site.

Word Processing (and More) with KWord

Long before OpenOffice.org was even a gleam in the eye of Tux, the Linux penguin mascot, Linux users still needed to create documents graphically. Over the years, a fine selection of quality business/office applications have been developed for both the GNOME and KDE desktop environments, including suites for each environment. For GNOME, the GNOME Office suite (`http://live.gnome.org/GnomeOffice`) is a loosely coupled grouping of business/office applications that includes applications such as the AbiWord word processor, the Gnumeric spreadsheet (discussed in Chapter 17), and a few other applications. By loosely coupled, I mean that there is little automatic interaction or data-sharing between the programs in the GNOME Office suite—it's really more of a fraternal organization than an integrated office suite.

On the KDE side, however, the story is very different. KDE's KOffice suite is a tightly coupled set of business/office applications that easily and automatically share data whenever possible. KOffice includes applications such as the KSpread spreadsheet (discussed in Chapter 17), the KPresenter presentation program (also discussed in Chapter 17), the Kivio flowcharting program (you guessed it, also discussed in Chapter 17), a few other applications, and the subject of this section, the KWord word processing application. (For more information about KOffice, see their web site at `www.koffice.org`.)

KWord is a powerful and popular word processing application with some interesting twists, such as support for doing layout. Traditional word processing applications, such as Microsoft Word, Corel WordPerfect, OpenOffice.org Writer, and so on, provide you with a blank page when you create a new document. You begin entering text, and that text naturally flows from one page to the next as you become more verbose. You can change the width of the text and the margins on the page, and can flow text around embedded graphics, but your document consists of a single monolithic sequence of text from the beginning of the document to the end of the document. Layout-oriented applications, such as QuarkXPress, Adobe InDesign, FrameMaker, and PageMaker, Scribus (discussed in the next section), and KWord use a much more flexible

mechanism known as *frames* to provide very detailed control over where text, illustrations, tables, and other portions of your documents are located on the pages of your document. Frames are rectangular areas that are created on the pages of your document, and in which any document element can be located. In KWord, frames can be grouped into *framesets*, which are a series of frames that contain related content. Text can flow from one frame to another in a frameset, which makes it easy to create multi-column documents or to do the standard newspaper trick of "this story continued on page N." (You remember newspapers, right?) Frames are also commonly used to group pictures and captions and simplify moving them around on a page. When you first create a document in KWord, you are presented with a blank page on which you can begin entering text, just like any other word processor. You can enter text on this blank page, but can also draw frames there, define how text in the basic document flows around the frames, define how text flows from one frame to another on subsequent pages, and so on. Frames provide a powerful model for creating visually appealing and well-designed documents, and KWord is unique among Linux word processors in its support for both frames and the standard page-oriented word processing model.

KWord isn't installed on Kubuntu systems by default, largely (I presume) because of the ubiquity of the OpenOffice.org suite as a reflection of the convergence of Linux desktop systems on a common set of applications. That's everyone's loss, because KWord is well worth a look, and any rumors of its demise are greatly exaggerated. KWord is not only actively supported on standard KDE 3-based Kubuntu systems, but is undergoing a major re-write/update to take advantage of the capabilities of KDE 4 (shown in a screenshot later in this section). On Kubuntu systems, installing and using KWord is a no-brainer because all of its dependencies are already installed, but if you have to create any document that is more complex than a simple series of pages, installing and using KWord is still worth considering on a standard GNOME-based Ubuntu desktop system.

NOTE **Because KWord is traditionally used on KDE desktops, the screenshots in this section were all done on a standard Kubuntu system.**

Installing KWord

To install KWord on a Kubuntu system, start the Adept Manager from the K ➪ System menu, and enter `kword` in the search area. Right-click on the `kword` package and select Request Install from the pop-up menu. This will also select the `kword-data` package for installation. Next, click Apply Changes in the Adept Manager, and the installation begins. Once the installation completes, exit the package manager, and you're ready to begin experimenting with KWord.

TIP **You can also install KWord on a standard GNOME-based Ubuntu desktop system using the Synaptic Package Manager. See Chapter 23, "Adding, Removing, and Updating Software," for more detailed information on installing applications using the Synaptic Package Manager.**

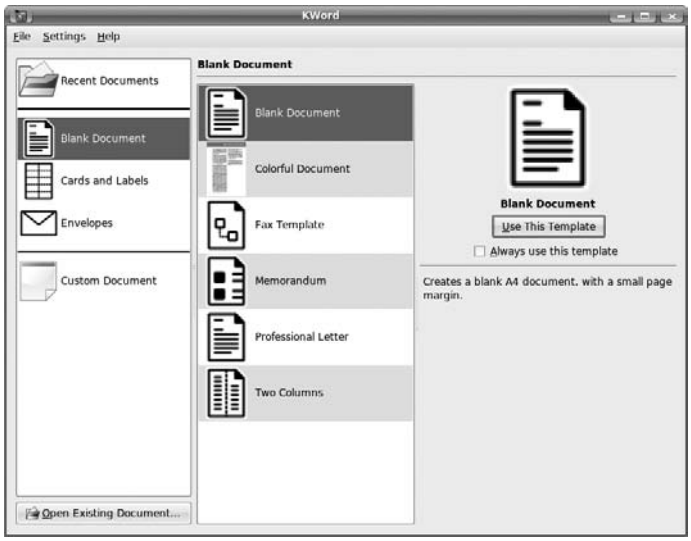
Quick Tour of KWord

You can start KWord by selecting the Office ➪ KWord-Word Processing menu item, or by executing the `kword` command in a Konsole, xterm, or GNOME Terminal window. When started from the menus or from the command line with no filename as an argument, KWord initially displays the screen shown in Figure 16-27, which enables you to create a new document based on any existing template by selecting that style and clicking Use This Template, or to open an existing document by clicking Open Existing Document and navigating to the document that you want to open.

You can click the Always Use This Template button to skip the template selection dialog in the future. Selecting this checkbox and a specific template will cause KWord to always display a blank document of the selected type when it starts. To re-enable the style selection dialog at startup, select File ➪ New and deselect the Always use this template checkbox.

FIGURE 16-27

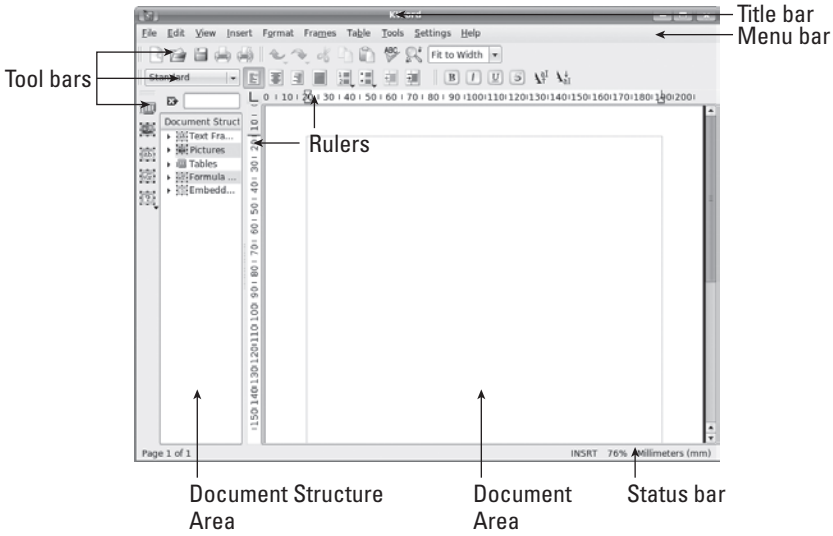
KWord's initial template selection dialog



When creating a new document (in this case using the Blank Document template), KWord displays an empty document in a screen like the one shown in Figure 16-28.

FIGURE 16-28

An empty document in KWord



As shown in Figure 16-28, KWord document windows are divided into the following primary areas:

- **Title bar:** Displays the standard KDE window decorations and the name of the document that you are editing (if any). When creating a new document, no document name is shown until the document has been saved.
- **Menu bar:** Displays the drop-down menus provided by KWord.
- **Horizontal and vertical toolbars:** Display related sets of commands as icons and related graphical input areas. By default, KWord displays the following toolbars:
 - **File:** Provides alternatives to the menu commands for common file operations such as opening and saving files.
 - **Edit:** Provides alternatives to the menu commands for common editing operations such as cutting, copying, and pasting.
 - **Insert:** A vertical toolbar (by default) that enables you to create and modify different types of frames.
 - **Paragraph:** Provides quick access to a drop-down–style menu and icons for various paragraph alignment and quick style commands, such as for creating bulleted and numbered lists.
 - **Format:** Enables you to quickly change the formatting of a selected portion of the document.

 You can modify the set of toolbars that are currently displayed by right-clicking on any toolbar to display the pop-up Toolbar menu and selecting or de-selecting any of the toolbars from the menu. I often disable the File and Edit toolbars because I tend to use menu commands or keyboard shortcuts for these operations and thus those toolbars are a waste of space for me. You can also modify the icons that are displayed on various toolbars by selecting the Configure Toolbars command from the Toolbar menu.
- **Document structure area:** Enables you to create, select, and modify the different types of frames and other structural document elements that KWord supports. These are text frames, pictures, tables, formula frames, and embedded objects. You can toggle whether the document structure area is displayed by selecting Hide Doc Structure/Show Doc Structure from the View menu.
- **Document area with rulers and scrollbars:** The working area for a new or existing document. By default, all documents are surrounded by horizontal and vertical rulers to help you accurately position document elements and to generally identify the current cursor position on a page. All documents also have horizontal and vertical scrollbars to help you navigate through your document. The horizontal scrollbar is useful only if the width of your document exceeds the width of the document area, such as when you have zoomed in on a document.

NOTE **Documents are displayed using a page view by default. You can select the View ➪ Display Mode ➪ Text Mode menu command to switch to an un-paginated view that displays your document as a single flow of text. Text mode will not display any frames that you have created.**

- **Status bar:** The left portion of the status bar displays the current page number and total number of pages. The right portion of the status bar displays the current text mode (either INSRT, to insert text at the current cursor position, or OVER, to over-write existing text beginning at the current cursor position), the current page-scaling factor, and the measurement units that are currently being used in KWord. The center portion of the status bar displays name, size, and coordinate information about the frameset with which any selected frame is associated.

KWord provides extensive online help that you can view at any time by pressing F1 or by selecting the Help ➪ KWord Handbook command. Unfortunately, the KWord documentation is currently out of date, especially regarding the use of frames and the old mechanism of selecting Page Layout templates. KWord used to differentiate between documents that used a single, full-page frame (like traditional word processors) and documents that used multiple frames. This difference has been eliminated, but the documentation hasn't quite caught up. You can generally ignore any references to the Page Layout templates or their associated button.

Creating and Editing Documents in KWord

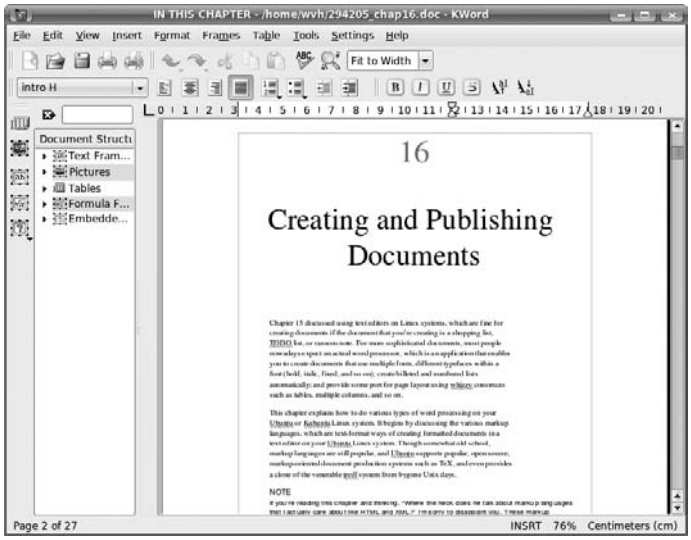
As discussed in the previous section, starting KWord initially displays a template selection dialog shown in Figure 16-27, which enables you to create a new document based on any existing style by selecting that style and clicking Use This Template, or to open an existing document by clicking Open Existing Document and navigating to the document that you want to open.

When opening an existing document from this dialog, KWord initially displays the contents of the Documents directory in your home directory. After selecting a file, click OK to open that file in KWord. If the file was produced by another word processor but is in a format that KWord understands, KWord will open the file without requiring that you explicitly import the document.

Figure 16-29 shows the text of this chapter after having been opened in KWord.

FIGURE 16-29

An existing document in KWord

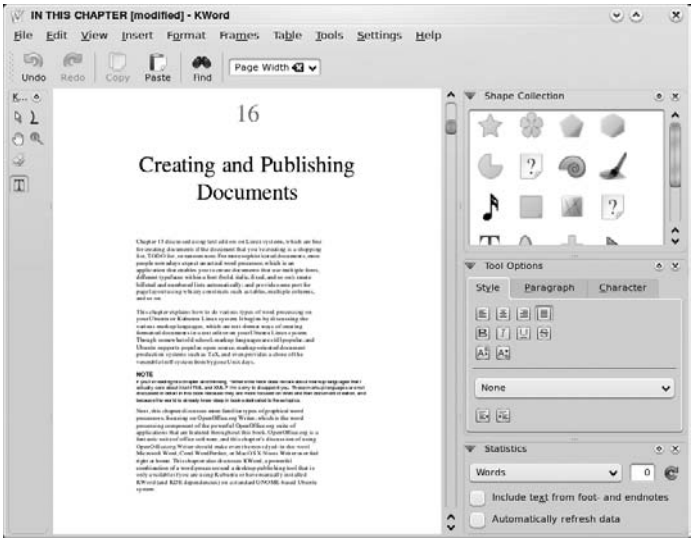


The discussion of KWord in this section focuses on KWord 1.*x*, which is the standard KDE 3.*x* version of KWord that is provided in the Ubuntu repositories for Ubuntu and Kubuntu systems. (The exact version used in this section is version KWord 1.6.3.) KWord is being substantially rewritten for KDE 4, and

will have a very different look and feel and a substantially different menu structure. Figure 16-30 shows the same document being edited in an alpha version of KWord for KDE 4.*x* systems. Something to look forward to!

FIGURE 16-30

Sneak Preview: An existing document in KWord for KDE 4



Back to KWord for KDE 3. Once you have opened or created a document in KWord, you can begin adding text to or editing that document just as you would in any word processor. You can enter text, select portions of the text or text blocks such as paragraphs and apply and modify styles by selecting them from the Paragraph toolbar, cut and paste text to move it around, and so on.

Working with Frames

KWord is an extremely usable and flexible word processor when all of the pages in your document contain full-page text frames, which is the default for documents that are created from the Blank Document template. However, adding other frames and controlling how text flows between them is the key differentiator between KWord and other word processors, and shows the real power of KWord.

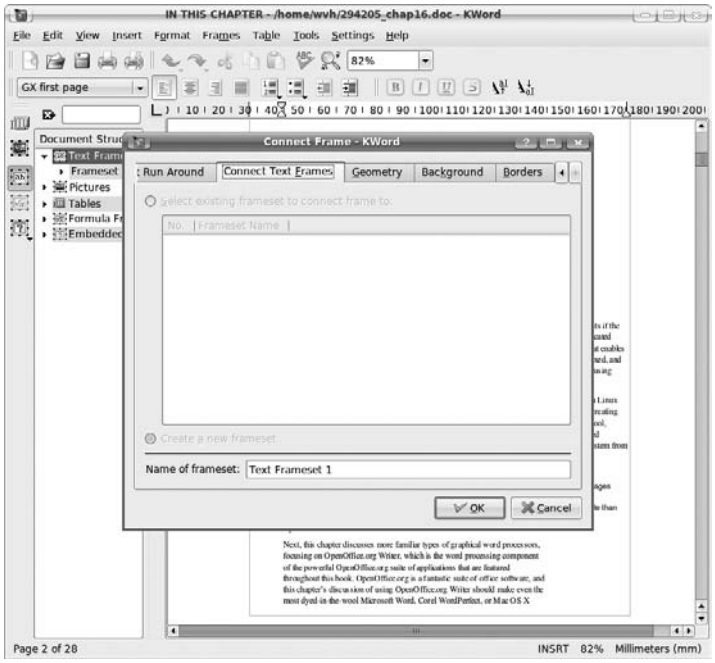
To add a frame to a document, you can do any of the following:

- Press F10.
- Select the Insert ➪ Text Frame menu command.
- Click the Create a new text frame icon in the vertical Insert toolbar that is shown by default on the left of a KWord window.

After doing any of these, position the cursor in the document area where you want one corner of the frame to appear, hold down the left mouse button, and drag the cursor to create the frame on the page. As you move the cursor, a rubber-band box (a box whose perimeter is composed of dashed lines and which stretches from the point where you first clicked the mouse button) displays to show an outline of the frame that you are creating. When you release the cursor, a dialog that enables you to define the properties of that frame displays over the current document window, as shown in Figure 16-31. This dialog is known as the Frame Properties dialog.

FIGURE 16-31

KWord's Frame Properties dialog



The default text flow used on the pages of a KWord document that was created using the Blank Document template is usually called Frameset 1 or Text Frameset 1. When you add a new frame, KWord assumes that you want to create a new text flow for that frame, and provides a new default name such as Text Frameset 1 or Text Frameset 2 (depending upon whether Text Frameset 1 already exists). You can see the new name at the bottom of the Frame Properties dialog's Connect Text Frame tab in Figure 16-31.

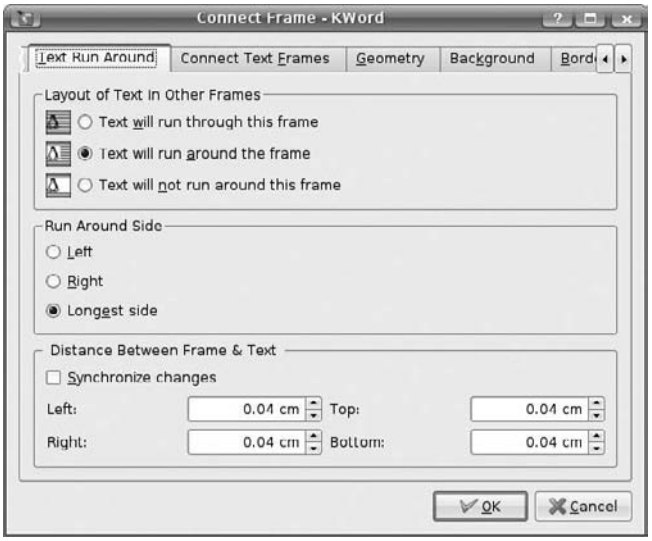
TIP **You can redisplay the Frame Properties dialog for any frame at any time by clicking on the border of the frame to select it, right-clicking on the border of the selected frame, and selecting Frame Properties from the pop-up menu.**

KWord provides a good set of default attributes for new frames, all of which you can see or modify on the various tabs of the Frame Properties dialog. For example, the Text Run Around tab, shown in

Figure 16-32, assumes that text in any frame that is beneath this one should run (wrap) around the new frame.

FIGURE 16-32

Text Run Around options for frames in KWord



You rarely have to modify the default settings when creating your first frame, so click OK to accept the default settings and close the Frame Properties dialog. Figure 16-34 shows a sample document with a new, empty frame, embedded on the page.

Similarly, the default settings on the Frame Properties dialog's Options tab, shown in Figure 16-33, prevent a new frame from being duplicated on subsequent pages unless explicitly requested, and also helps ensure that no text is lost if the text for a frame is too long for that frame, by automatically increasing the size of the frame to hold that text.

To demonstrate how to chain frames together by adding new frames to existing framesets, let's first enter more text in the new frame than it can hold, based on its size. (This is just an example—not a suggestion for how to lay out your documents!) Figure 16-35 shows the frame after adding "too much" text—the frame has automatically resized to accommodate the new text.

Let's now add another frame to the same page. Click the Create a new text frame icon in the vertical Insert toolbar, and drag and draw the outline of the new frame on your page. The Frame Properties box displays, as shown in Figure 16-36.

In this case, note that the Frameset that you created when you added the first frame is now listed in the frameset list. To add your new frame to the same frameset as the first frame that we created, select the Select existing frameset radio button, select the frameset that you want to attach the new frame to, and click OK. Your screen will now look something like the one shown in Figure 16-37, depending on where you added the frame.

FIGURE 16-33

Basic options for frames in KWord

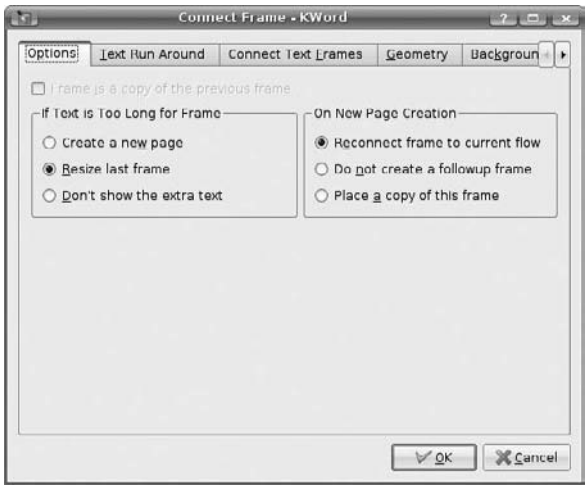


FIGURE 16-34

New frame in an existing document

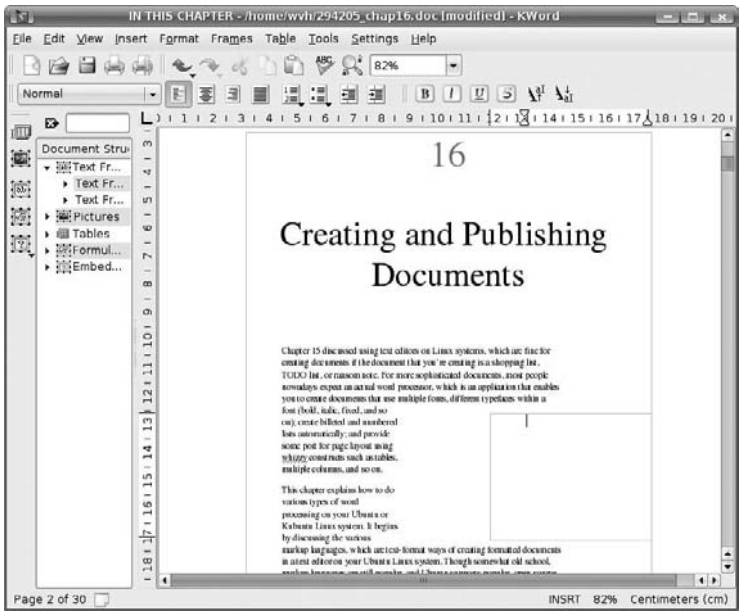


FIGURE 16-35

Automatic frame expansion for new content

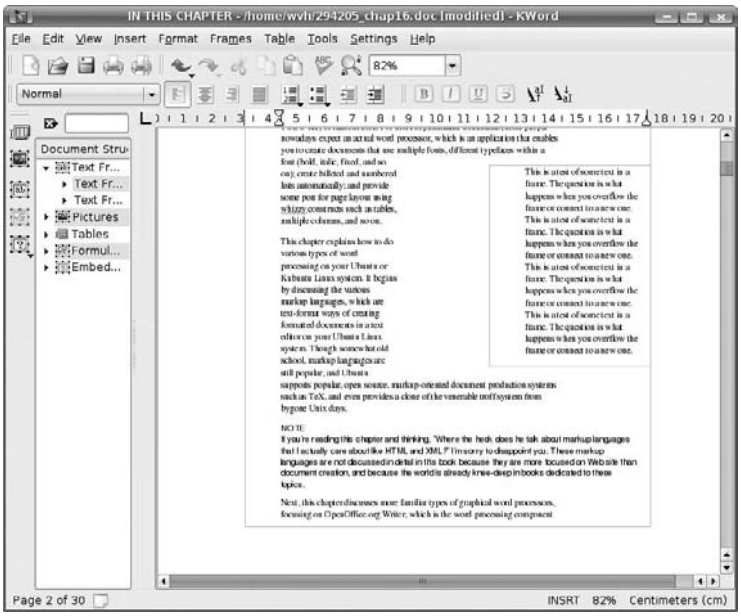


FIGURE 16-36

The Frame Properties dialog for another frame

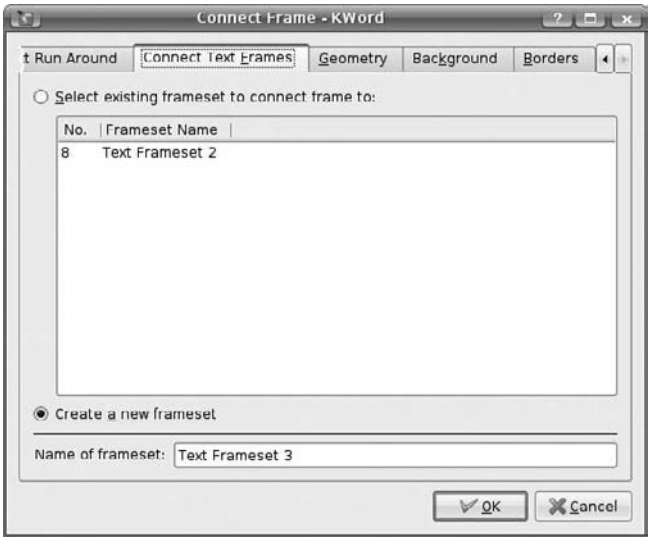
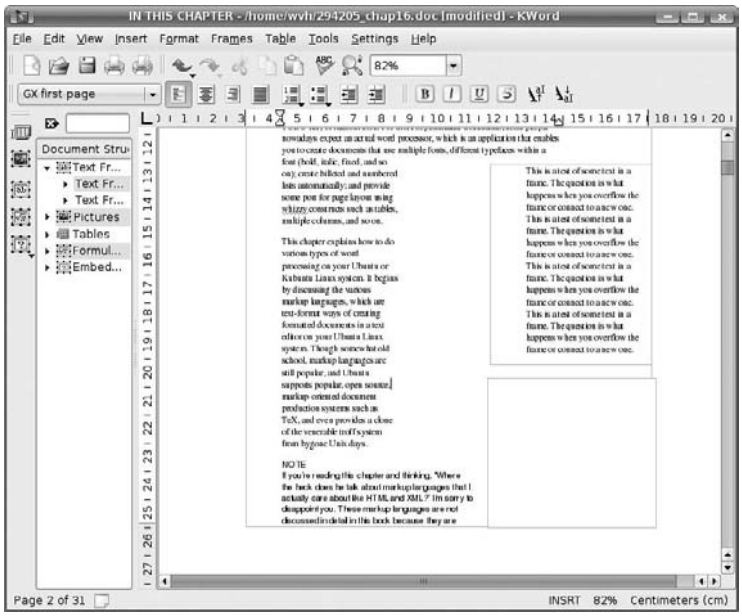


FIGURE 16-37

Two frames on the same document page



Note that the first frame did not automatically return to its initial size. To resize it, click the border of the top frame, select the handle in the middle of the bottom border, and drag it to resize the frame. As you do so, you will see any text that does not fit in the top frame automatically flow to the bottom one, as shown in Figure 16-38. The content on the text page will also adjust itself to reflow around the modified frame boundaries.

The whitespace in and around frames can be adjusted in different ways:

- By displaying the Frame Properties dialog's Text Run Around tab for that frame and adjusting the Distance Between Frame & Text values
- By adjusting the indentation and margins of the text elements that wrap around the frames
- By adjusting the indentation and margins of the text elements within the frames

Figure 16-39 shows the same document as that shown in Figure 16-38, but with the paragraph margins and indentation adjusted for the paragraphs within and around the frames. Looks much better, I think.

FIGURE 16-38

Two frames after resizing

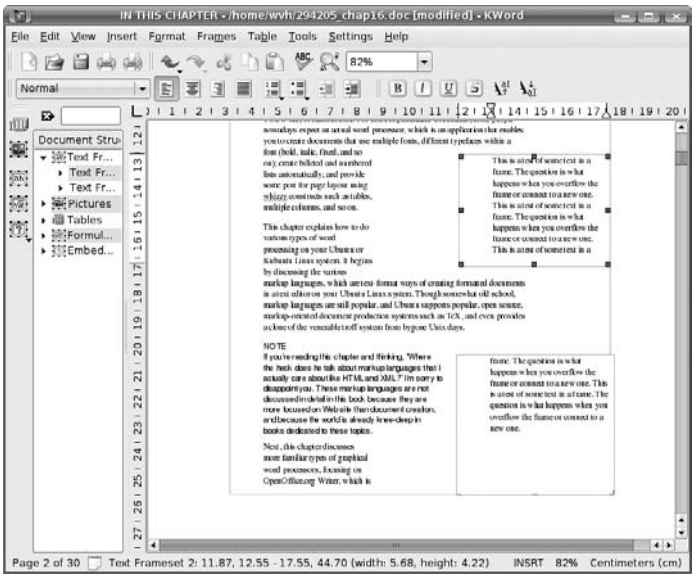
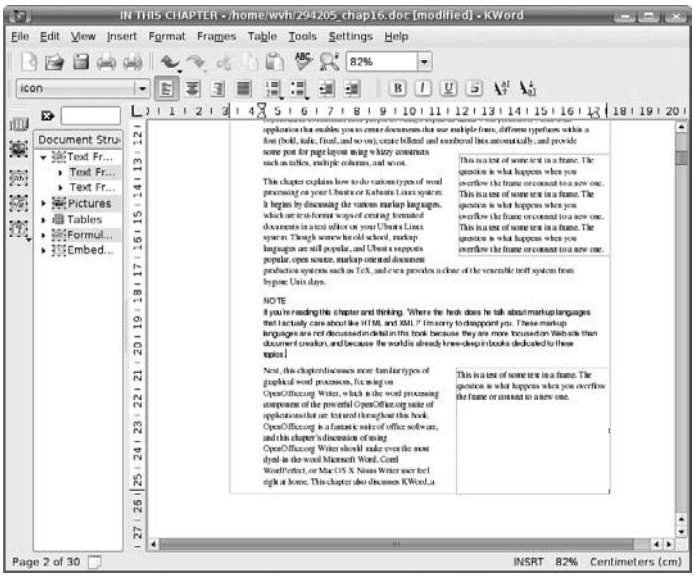


FIGURE 16-39

Frames and surrounding text after adjusting margins



Importing, Exporting, and Saving Documents

Saving the document that you are editing in KWord works the same way that it does in most other word processors, by using the File menu's Save command. If the document is already associated with a file (such as if it has been saved before), the file is saved without any subsequent messages. However, if the document does not already have a name (such as when you have not saved it before), a Save As dialog displays. This dialog enables you to navigate to the directory where you want to save the file, supply a name for the file, and specify the Filter used to save your file. This filter defines the format in which your file is saved.

NOTE **When saving or exporting documents that use KWord-specific features such as frames, a dialog displays if the target document format that you specify does not support all of the features that your document uses. The KWord output format is always a safe choice if you will be continuing to work with this file in KWord.**

The KWord folks understand that you may have started out with some other word processor, and therefore have some legacy documents that you will want to continue to use KWord to work with. To solve this problem, KWord provides a File ➪ Import command that enables you to import documents in a number of supported formats so that you can begin working with them in KWord. Similarly, the developers realize that not everyone has KWord, and that you may need to produce documents in other formats for easy interchange with your peers. To address this problem, KWord provides a File ➪ Export command. For a complete list of the import and export formats supported by KWord 1.6, see `http://koffice.org/filters/1.6`. Suffice it to say that most common formats, such as Microsoft Word, Corel WordPerfect, OpenOffice.org, RTF, SGML, and XML generally work quite well.

On the other hand, the list at `http://koffice.org/filters/1.6` can be somewhat misleading. For example, the SGML and XML export formats export every document element as a paragraph—which is syntactically valid, but not necessarily accurate (unless your SGML and XML documents are very different from mine).

More Information About KWord

This section has only scratched the surface of the features, menu commands, and dialogs provided by KWord. For more information about KWord, see the following links:

- **Home Page:** `www.koffice.org/kword`
- **User Manual:** `http://docs.kde.org/development/en/koffice/kword/`
- **KOffice Supported File Formats:** `http://koffice.org/filters/1.6`

KWord is a well-known, stable, and powerful word processing application with some unique capabilities. Moving KWord to version 2.0 and porting it to KDE 4 should not only provide new features, but should ensure the continued use and popularity of KWord for the foreseeable future.

Desktop Publishing with Scribus

The previous sections of this chapter discussed how to create different types of documents on Ubuntu and Kubuntu systems using command-line and graphical word processors. These tools provide powerful styles to standardize the appearance of repeated elements in your documents. However, when the presentation and layout of your document is just as important as its content, you need a different sort of tool.

Word and text processors do just that, formatting and organizing the content in your documents, usually providing little support for actual page layout beyond controlling the page margins, size of the print area, and the number of columns of text that are displayed on each page. To do actual, detailed document layout, you need an application that combines word processing capabilities with the layout and detailed formatting capabilities required by the publishing industry—in other words, a desktop publishing package. It's a good thing that the open source folks ran into that problem, too, and an excellent (and free) desktop publishing package for Linux, named Scribus, is easily installed on your Ubuntu system.

Scribus has been under development for a few years now, and is an open source application that is available at no cost for Linux, Apple Mac OS X, and Microsoft Windows systems and uses the QT graphical interface so that it looks and behaves the same way on all supported systems. The Scribus home page is `www.scribus.net`—check there for the latest Scribus information.

As discussed in the introduction to the KWord section earlier in this chapter, the basic concept in desktop publishing is the frame, which is conceptually similar to the idea of a paragraph in a word processing system. A frame is the key document element in desktop publishing. Desktop publishing documents are composed of many different frames, which can contain text, graphics, and so on. Dividing documents into these compositional elements makes it easy to move or resize them independently to suit changes in the document, its target layout, and so on. Text frames can be linked so that text flows from one frame to another, but they are still separately manageable document elements.

The next few sections explain how to install Scribus, how to create documents using the templates that it provides, and what basic steps are necessary to create new documents with the detailed page layout required by your most demanding publishing work. It's not quite Quark XPress, Adobe InDesign, or Adobe PageMaker, but you can't beat the price and the quality is quite impressive. You can't import existing InDesign, PageMaker, or Quark files yet, although you can import OpenOffice and other open format documents as a starting point for your Scribus documents.

Installing Scribus

Because not everyone needs or wants a desktop publishing package, Scribus isn't installed by default on standard Ubuntu or Kubuntu Linux systems. The primary package that you need to install in order to use Scribus is the `scribus` package, but you may also want to install the `scribus-doc` and `scribus-template` packages to help you make the most of Scribus.

To install Scribus on a standard GNOME-based Ubuntu desktop system, start the Synaptic Package Manager from the System ➪ Administration menu, select Search, enter **Scribus** in the Search box that displays, and click Search. Once the search completes, right-click on the `scribus` package and select Mark for installation from the pop-up menu. You will be prompted to confirm that you also want to install the packages that it requires. Click OK. You should also select the `scribus-doc` and `scribus-template` packages for installation unless disk space is absolutely critical or you're already an experienced Scribus user.

CAUTION **To install Scribus on a Kubuntu system, start the Adept Manager from the K ➪ System menu, and enter** scribus **in the search area. Right-click on the** `scribus` **package and select Request Install from the pop-up menu. If you're interested in the documentation and extra templates (and you should be, unless you're already a hard-core Scribus user), you should also scroll down to where you see the** `scribus-doc` **and** `scribus-template` **packages, and request the installation of those packages, too.**

After selecting the package that you want to install, click Apply in the Synaptic toolbar (or Apply Changes in the Adept manager) to install Scribus on your system. Once the installation completes, you're ready to begin working with Scribus! Installing Scribus on your Ubuntu or Kubuntu system installs a Scribus menu item on the Office menu that you can use to run the application.

See Chapter 23, "Adding, Removing, and Updating Software," for more detailed information on installing applications using the Synaptic Package Manager or Adept Manager.

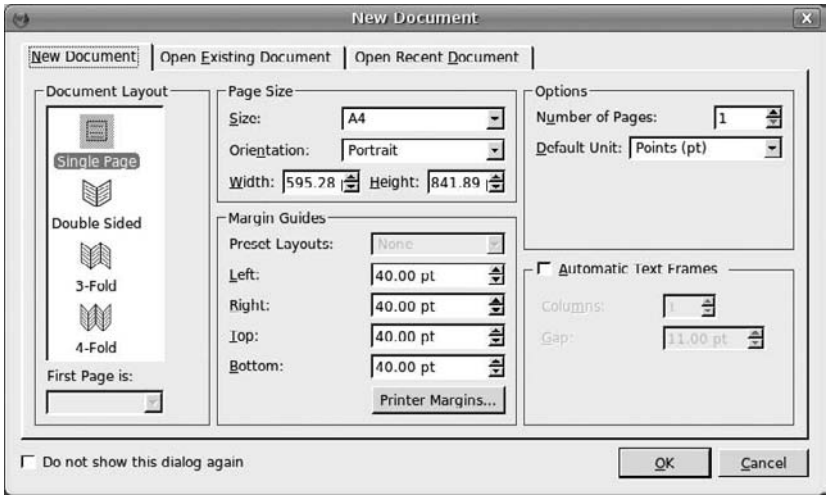
Taking a Quick Scribus Tutorial

Like OpenOffice.org Writer and most other publishing-related applications, Scribus comes complete with a few templates to get you started. When you first use a new application, it's often easiest to start with someone else's document, and then modify it to suit your purposes. Once you get the hang of using the application, you can create your own documents, as explained in the next section. This section provides a whirlwind tour of the key features of Scribus to get you comfortable with the basics.

To start Scribus, select Applications ➪ Office ➪ Scribus (or K ➪ Office ➪ Scribus on a Kubuntu system). Scribus takes a little bit of time to start up as it examines your system's font resources (which are critical to a detailed, layout-oriented application). The first time you start Scribus, you will see a new document dialog, as shown in Figure 16-40.

FIGURE 16-40

Selecting a template as a starting point



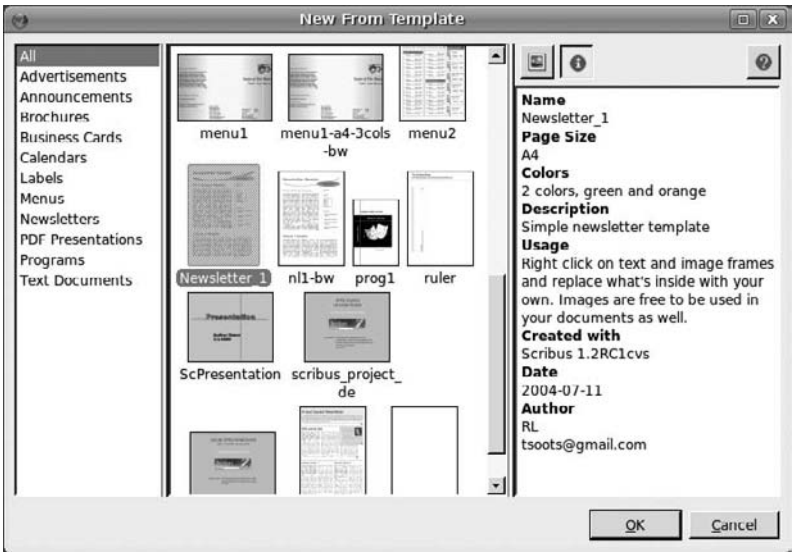
This dialog enables you to identify the document layout that you want to use when creating a new document. You can create a new document from a basic layout specification or using one of the templates provided with Scribus, as we'll do in this tutorial. To bypass this dialog for now, click Cancel. If you don't want to see this dialog the next time that you start Scribus, click the Do not show this dialog again checkbox before clicking Cancel.

Once Scribus starts, you can create a new document from one of its templates by doing the following:

1. Select the File ➪ New from Template menu command to display the dialog shown in Figure 16-41.

FIGURE 16-41

Selecting a template as a starting point



2. Select the template that you want to use for your new document and click OK. In this example, I'll use the Newsletter_1 template, which provides a working model for a newsletter that includes graphics. As Scribus opens the template, you may see one or more warning dialogs if the template uses fonts that are not installed on your system. For each warning dialog (if any), write down the font that was missing and click OK to substitute a similar font that is present on your system. You can subsequently use your system's package manager to search for and load any available font packages.
3. Once font processing completes, the new document that is based on the selected template displays in the main Scribus window, as shown in Figure 16-42.
4. Now things get interesting. To work with or simply change the text in any frame, first select that frame. You can then select the Edit Contents of Frame button in the toolbar (this is the icon that shows a white letter A and a text cursor). After selecting this command, you can directly add or delete text in the current frame. However, for more detailed changes involving font and text changes, I tend to use the Edit Text dialog. To access this dialog, you can simply click the Edit the text with the Story Editor icon in the toolbar (the icon that looks like a pen and tablet computer). I tend to access this through the context-sensitive menus rather than the toolbar, so you can also right-click on that frame to display the pop-up menu shown in Figure 16-43.

FIGURE 16-42

Scribus showing your working copy of a template

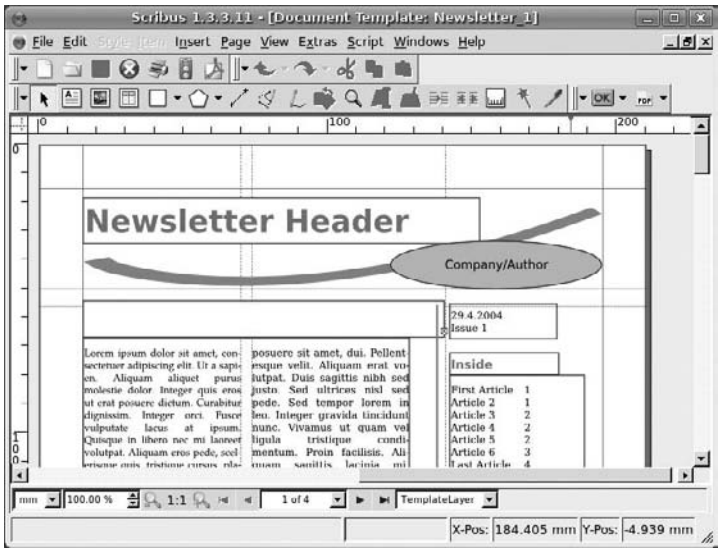
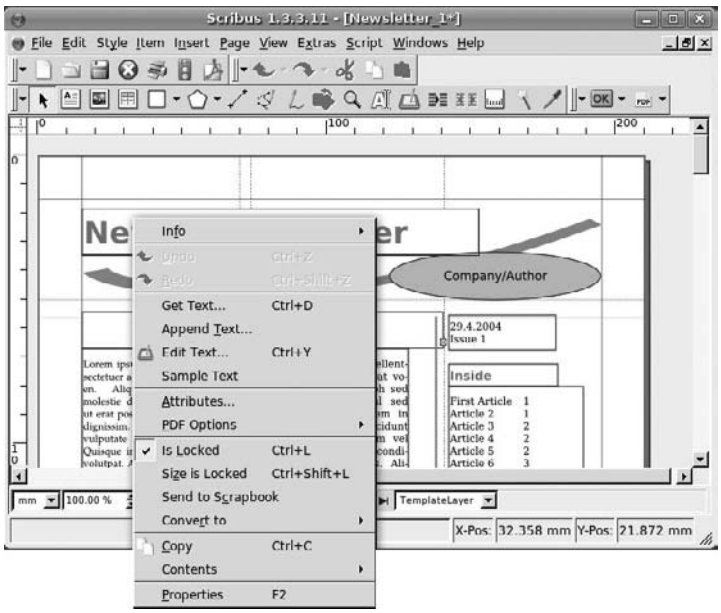


FIGURE 16-43

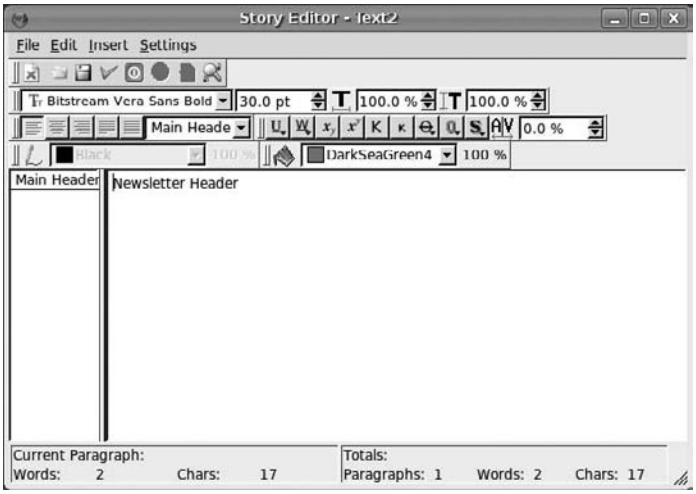
Editing options for frames in Scribus



5. To view the dialog used for low-level font changes, select the Edit Text menu item. A Story Editor dialog like the one shown in Figure 16-44 displays. The Story Editor is a small word processor, encapsulated within Scribus. Using the Story Editor, you can modify, replace, or add to the text in the frame using the lower window. The controls available in the Story Editor enable you to change justification within a frame, fonts, and font characteristics.

FIGURE 16-44

Modifying the text in a frame with the Story Editor



6. After modifying the frame text or its characteristics, click the checkmark icon or select the File ➪ Update Text Frame and Exit command. In this case, I've changed the title of the newsletter from "Newsletter Header' to "Linux People and Pets." An updated version of your document with the new information displays, as shown in Figure 16-45.
7. Working with graphical frames is similar but uses an external application (GIMP, which is discussed in Chapter 18, "Working with Graphics") to edit any existing image. For example, scroll down in Scribus until you see a figure on the right side of the second page. To replace this graphic, first right-click on its frame and de-select Is Locked from the context-sensitive menu to unlock the frame. This enables modifications to the frame. As you're working, you can lock any frame that you're happy with to prevent accidental modification.
8. To delete the existing image, right-click on the frame and select Contents ➪ Clear. The document updates to reflect the absence of the graphic, as shown in Figure 16-46.
9. To select a new graphic for the frame, right-click on the frame and select the Get Image menu command. The dialog shown in Figure 16-47 displays, enabling you to navigate to, select, and preview any existing graphic on your system.

FIGURE 16-45

The updated version of the document

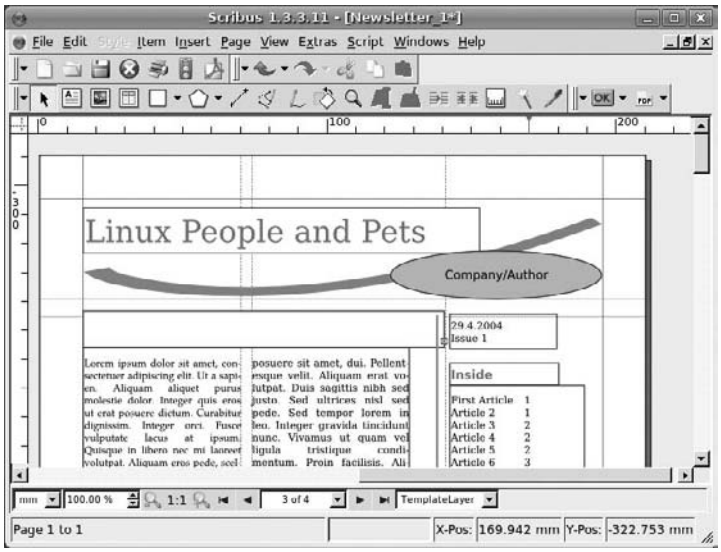


FIGURE 16-46

The newsletter after a graphic has been removed

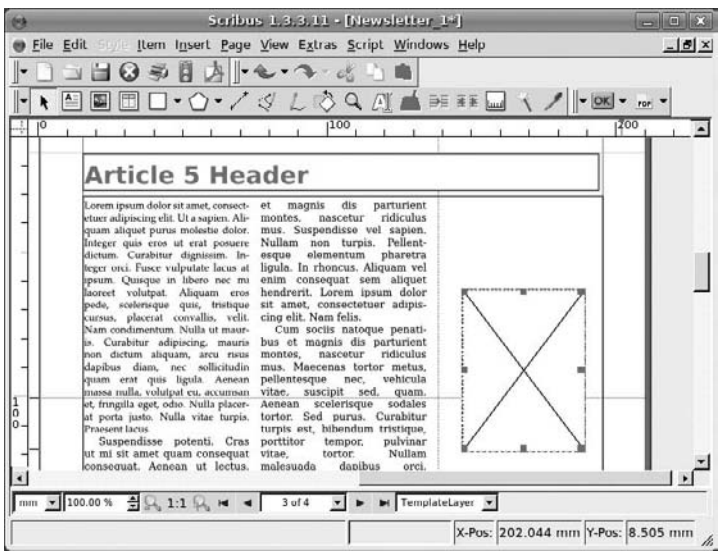
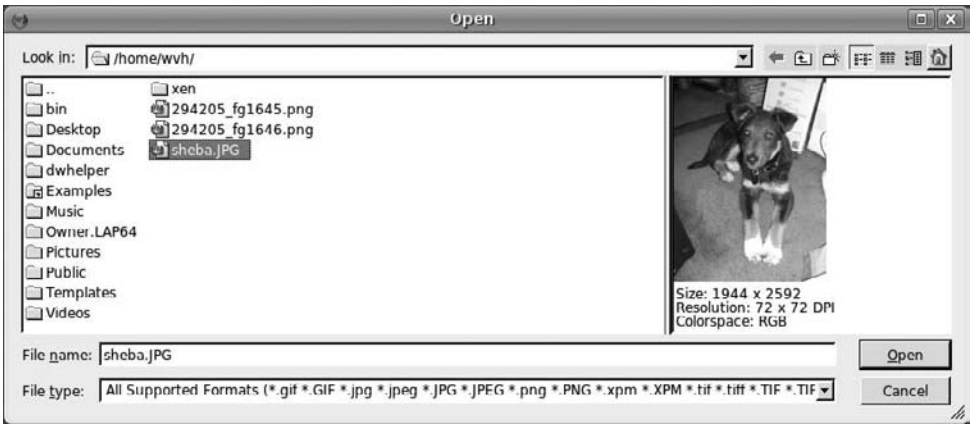


FIGURE 16-47

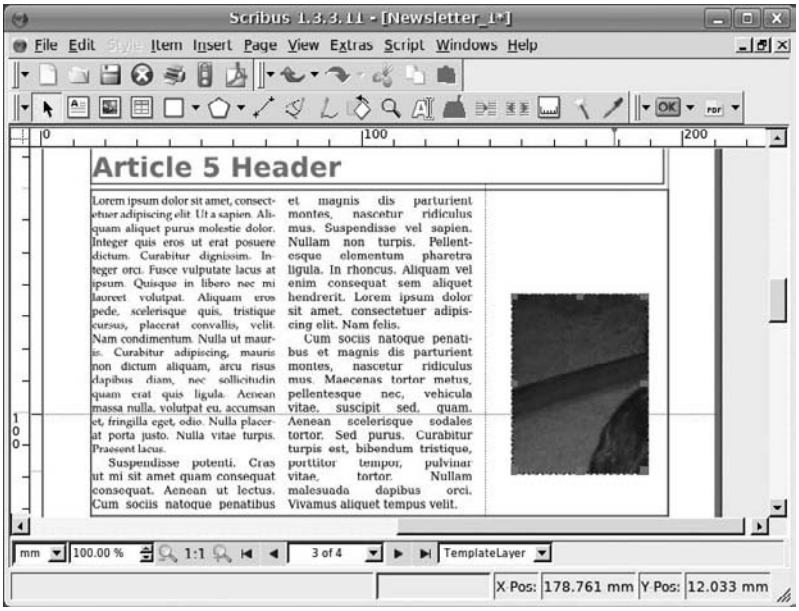
Selecting a new graphic for a frame



10. After selecting the new graphic, click Open to open the new graphic in the selected frame. The view of your newsletter updates to reflect the new graphic, as shown in Figure 16-48. Hmm, that doesn't actually look very good because the figure is supposed to show a picture of my dog.

FIGURE 16-48

The updated document showing the new graphic

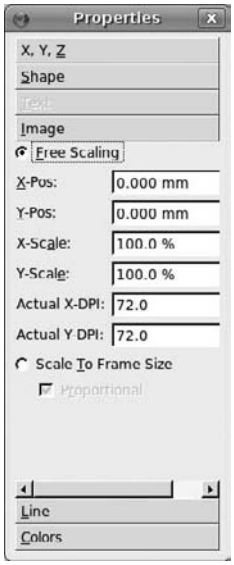


11. Because the graphic is a different size than the original frame, right-click the graphic and select the Properties menu command to display the Properties dialog, as shown in Figure 16-49. Click Image to display the Image Properties tab.

12. Select the Scale To Frame Size radio button to scale the image to fit the frame. To completely fill the frame, de-select the Proportional setting, which would otherwise preserve the original X and Y ratio of the image. Click the close box in the upper-right corner of the Properties dialog to close it. Figure 16-50 shows the scaled image in the frame—much better!

FIGURE 16-49

The Properties dialog's Image tab



13. The extra space available in the document provides a good location for a new text frame. First, drag the image frame up by clicking and holding the left mouse button on the frame, and then dragging the frame to its new location. To add a new frame, click the Insert Text Frame icon from the toolbar (it looks like a small document) and click and drag to insert the new frame. For the purposes of this example, I'll insert some sample text into the frame by right-clicking the frame and selecting Sample Text from the pop-up menu. Select the language for the sample text from the dialog that displays, and click OK. It may seem odd to have a command that inserts random text into a frame, but this is useful when laying out a document before the final content is ready, and it saves you the trouble of having to locate some random text, copy it to the Clipboard, and then paste it. Your document updates to reflect the new frame and its contents, as shown in Figure 16-51.

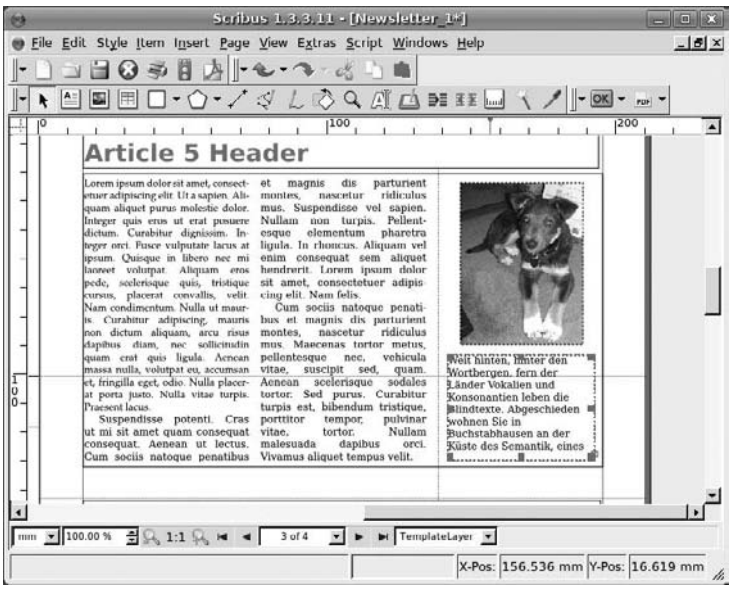
FIGURE 16-50

The scaled image in the document



FIGURE 16-51

The updated document showing the new text frame

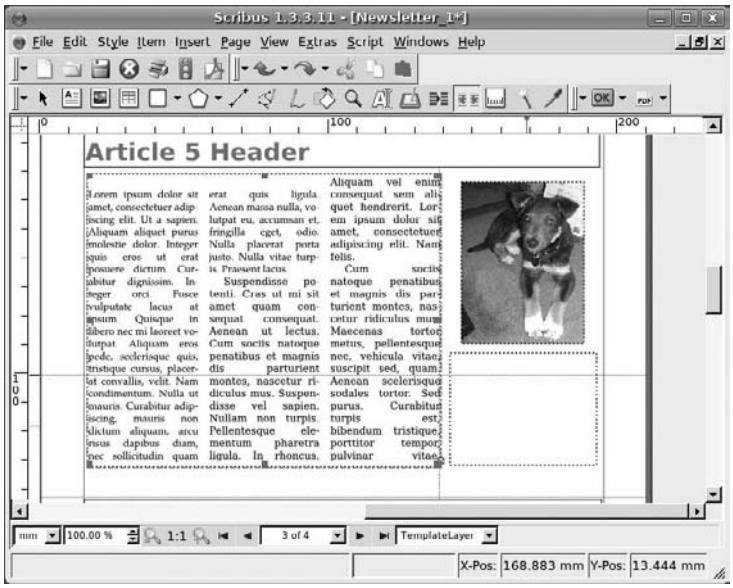


14. The final step in this quick tutorial is to link two text frames so that any text that overflows from the first goes into the second. You can link any two frames, but the frame that you are linking to must be empty. To clear the frame that you just added, right-click on it and select Content ➪ Clear. (When clearing existing frames, you may first have to unlock their contents by de-selecting Is Locked from the content-sensitive menu.)

15. Resize the frame to the left of the figure and frame that we just added. First, right-click on the frame and de-select Is Locked. Red handles appear on the outline of the frame. Select the one in the middle of the right size of the frame, hold down the left mouse button, and drag the frame border to the right. The document should look something like that shown in Figure 16-52.

FIGURE 16-52

The updated document resizing a text frame



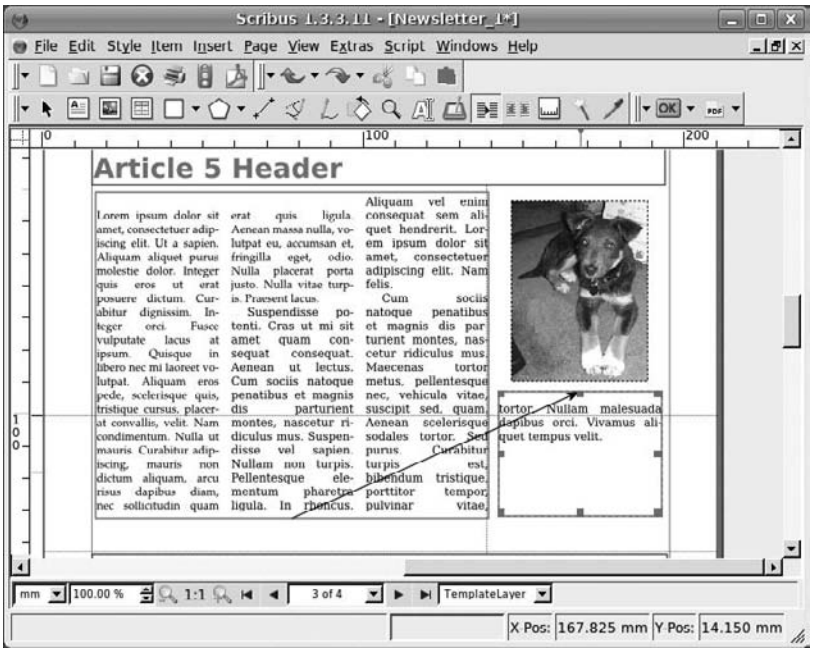
16. Because the frame is still selected, click Link Text Frames button in the toolbar. (It shows two text frames with an arrow beside them, and is to the right of the Edit the text with the Story Editor button.) Left-click on the frame that you want to flow text to, which is the frame that we added earlier. An arrow displays on the screen, showing that the two paragraphs are now linked. The screen updates, flowing text from one frame to the other, as shown in Figure 16-53.

That's it for this tutorial. You can now save your document, or export it for printing by selecting the Save the current document as PDF icon (identified by the familiar PDF graphic).

Scribus is a powerful application with many features, and deserves its own book some day. This tutorial guided you through the basics of creating, modifying, and linking text frames, as well as inserting and modifying graphics. Scribus is an impressive program that is a shining example of the power of the open source model and the skill and determination of a few developers who saw a need, and filled it.

FIGURE 16-53

The updated document showing linked text frames



More Information About Scribus

Scribus is a relatively young application, but a significant and ever-increasing amount of information about it is already available on the Web, in various forums, and in different parts of the Scribus Web site. My favorite locations for HOWTO information about Scribus are as follows:

- Scribus Documentation Site (`http://docs.scribus.net/`) provides a good deal of information about setting up and using Scribus, printing, color management in Scribus documents, and so on.
- Scribus Wiki (`http://wiki.scribus.net/`) provides a great forum for discussing and documenting how to use Scribus.
- Tutorial: Getting Started with Scribus (`http://docs.scribus.net/index.php?lang=en&page=tutorials/freedomyug/scribustutorial`), although written for an older version of Scribus and thus slightly dated, is an excellent hands-on tutorial for getting started with Scribus. The sample files used in the tutorial are also available for download from this site.

Scribus is an open source project staffed by volunteers. If you become a Scribus fan and have something to contribute, I'm sure they'd be glad to hear from you.

Other Word Processors and Office Suites

This chapter focused on the most popular, open source command-line and graphical applications for word processing and page layout. There are certainly others—this is open source, after all, where anything worth doing is worth doing several times because of different perspectives on the task at hand. The following are two other popular word processors for Linux:

- **AbiWord:** The AbiWord word processor is part of the GNOME Office Suite (`live.gnome.org/GnomeOffice`). AbiWord is a powerful and popular word processor for Linux, with excellent capabilities for importing and working with existing documents in formats such as Microsoft Word. AbiWord is available in the Ubuntu repositories, and can easily be installed on your Ubuntu or Kubuntu system using the appropriate package manager for your system.
- **Applixware:** Applixware is an office suite that includes a powerful word processor, but has had an on-again, off-again relationship with Linux. Originally marked as an end-user package that was a direct competitor to StarDivision's original StarOffice Office Suite, Applixware has morphed into more of a toolkit for developing real-time and analytical applications. For more information about Applixware, see `www.vistasource.com/vs2/en/index.php`. For an article on installing Applixware on Ubuntu Linux, visit the Applixware group at `http://groups.google.com/group/Applixware` and search for Ubuntu.

Both of these word processors are quite easy to use and have their own circle of devotees. If you find something that you don't like about OpenOffice.org Writer or KWord, installing and experimenting with any of these is well worth your time. After all, you can't beat the price.

Summary

This chapter explained how to do several word processing tasks on your Ubuntu system. It began by discussing TeX and LaTeX, the classic, command-line formatted document processors for Linux and UNIX systems, and even highlighted some of the graphical tools that are available for working with LaTeX documents. Next, the chapter covered OpenOffice.org Writer, the word processing component of the powerful OpenOffice.org office suite, and explained how to create documents and templates in Writer. This was followed by a discussion of KWord, a powerful word processing and layout application that is typically used only on KDE-based systems, such as Kubuntu, but which can also be installed and used on standard GNOME-based Ubuntu desktop systems. The chapter also covered Scribus, a powerful desktop publishing application for Linux, which provides the layout and formatting power you need when appearance and detailed layout are just as important as content. The final section of the chapter listed two other word processing packages that were not discussed in this chapter but which are also available for Ubuntu.

The document production software packages provided with Ubuntu Linux are great, but people do not live by documents alone. Chapter 17 explains how to use other office-related software on your Ubuntu system to create other popular types of files, such as spreadsheets and presentations.

Chapter 17

Other Office Software: Spreadsheets and Presentations

The previous chapter discussed how to create documents on your Ubuntu and Kubuntu systems. Along with surfing the Web and sending and receiving e-mail, writing letters, keeping To Do lists, and working on your novel are probably the most common personal tasks that you'll find yourself doing on your new machine. Word processing is also a fairly critical business task because, although most companies live on e-mail nowadays, there's still always the need for the occasional memo, planning document, and so on.

This chapter discusses the other commonly used and office-related software that's available for your Ubuntu and Kubuntu Linux systems, namely spreadsheets and presentation software. While few people actually need to give presentations on their home computer or to their family (maybe "Next Month's Budget — What It Means to You"?), spreadsheets are a crossover task that is often as important to home users as it is to office workers. If you are working from home or using Ubuntu or Kubuntu on a laptop for business reasons, you may need to create both spreadsheets and presentations, and both Ubuntu and Kubuntu make this easy (well, as easy as it is on any computer system). As you might expect, Ubuntu and Kubuntu systems install slightly different spreadsheet and presentation packages by default:

- Ubuntu Linux Desktop, OEM, or Text-Mode installations automatically install the OpenOffice.org Calc spreadsheet application as part of the OpenOffice.org suite, and the easy-to-use Synaptic Package Manager makes it easy to install other spreadsheets (such as the lighter-weight but equally powerful Gnumeric spreadsheet package).
- Ubuntu Linux Desktop, OEM, or Text-Mode installations also automatically install the OpenOffice.org Impress presentation application as part of the OpenOffice.org suite.

IN THIS CHAPTER

Using spreadsheets: A tutorial

Using OpenOffice Calc

Using Gnumeric

Using KSpread

Creating presentations with Impress

Creating presentations with KPresenter

- Kubuntu automatically installs the OpenOffice.org Calc spreadsheet application as part of the OpenOffice.org suite, and also installs the KSpread spreadsheet packages, which is part of the Koffice suite.
- Kubuntu also automatically installs the OpenOffice.org Impress presentation application as part of the OpenOffice.org suite. The Adept package manager makes it easy for you to install other, lighter-weight presentation software, such as KPresenter, which is part of the Koffice suite.

NOTE **To save repeating the phrases "OpenOffice.org Calc" and "OpenOffice.org Impress" (or OpenOffice.org Presentation, as it's known on the Ubuntu menus) more times than necessary throughout the rest of this section, I'll typically refer to OpenOffice.org Calc simply as Calc (its name among friends) and OpenOffice.org Impress (i.e., Presentation) simply as Impress.**

Introduction to Spreadsheets: A Quick Tutorial

This section provides background information on what a spreadsheet is, how it works, and how to use one. If you're already familiar with spreadsheets, you can skip this section. However, if you're not, this section is for you.

In paper terms, a spreadsheet is a sheet of paper divided into rows and columns that break down financial transactions into separate entities that a business person can examine and analyze. The first computerized spreadsheet was created for a mainframe computer system in the early 1960s by a professor named Richard Mattessich, then at the University of California at Berkeley. As a mainframe application, this had a somewhat limited audience — for you and me, spreadsheets were one of the original killer applications for personal computers, going all the way back to VisiCalc on misty dawn-of-time systems such as the Apple II. Then, like now, the power of the personal computer gave people the ability to enter and manipulate financial data on their own systems. Spreadsheets not only are good for calculations related to actual expenses and other financial transactions, but provide an excellent mechanism for "what if" calculations. What if we bought a car now — could we afford the payments based on our monthly income? What if I paid something extra on my mortgage — how long before it would actually be paid off?

In a computer spreadsheet, each rectangle in the spreadsheet is referred to as a *cell*. Each row on a spreadsheet typically provides a variety of information about a specific transaction or other entity, while the columns are used to visually compare, mathematically summarize, or perform some other mathematical operation on the value for each row that appears in that column. The following is an example in text form:

Meal	Cost	Tax	Tip
Breakfast	5.79	0.41	1.25
Lunch	8.48	0.60	1.75
Dinner	13.23	0.93	2.75

This simple table reflects my meals for a (fictitious) day. The first row is a set of headings that identify the meaning of the entries in each column. Each row provides the details of each of the meals for that day. It would be easy enough to find out the total cost for my meals that day in various ways — either by appending to the end a summary column that contains the total cost of each meal, and then adding all of the values in that column together, or by inserting a new row that contains the totals of all of that type of item for the day, and then adding all of the value in that row together. I prefer the first approach, a summary column, so

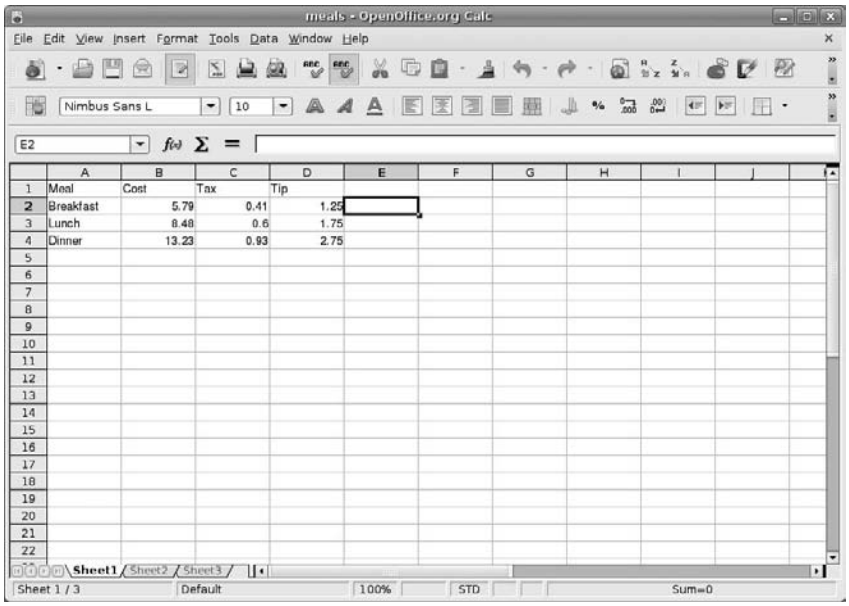
let's go with that approach for the remainder of this section. There's nothing preventing you from calculating your daily totals the other way if you prefer (or doing it both ways, for that matter) — the ability of spreadsheet applications to automatically perform these calculations is one of their coolest features.

NOTE **From this point on, this section is a hands-on tutorial for creating a simple spreadsheet. The instructions provided in the rest of this section work essentially the same way whether you're using OpenOffice.org Calc, KSpread, or Gnumeric. If you want to follow along, skip ahead to the section on starting your spreadsheet of choice, and then jump back here to work through the tutorial. It may be simpler to follow along in Calc because that is already installed on both Ubuntu and Kubuntu systems, but there are plenty of good reasons to use KSpread or to install and use Gnumeric. The screenshots in this tutorial show the example in Calc, but they will look exactly the same in KSpread or Gnumeric, with respect to differences in the menus, toolbars, and window decorations.**

Manually performing these calculations is fine if you don't have a computer or if you're far better at math than I am, but I prefer to use the appropriate tool whenever one is available — and spreadsheets come to the rescue in this case. I'd therefore start up a spreadsheet application (more about how to do that later), and enter the values from my paper table in the appropriate cells of the spreadsheet. (To enter a specific value in a cell in a spreadsheet, select that cell using the mouse or by using the arrow keys on your keyboard, and type the value.) Figure 17-1 shows the sample values after they've been entered into a sample spreadsheet.

FIGURE 17-1

Sample values in a Calc spreadsheet

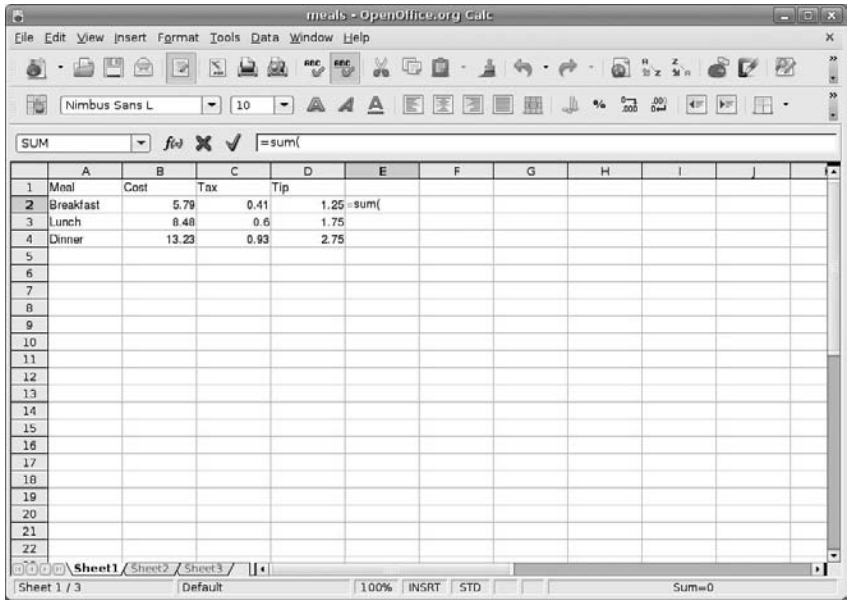


Now the fun begins! To create a cell that contains the sum of selected values from each row, click the cell where you want to create the value, and enter the value `=sum(`. The equals sign tells the spreadsheet that you are using a function, `sum` is the name of the function that you want to perform, and the opening parenthesis identifies the fact that this will be followed by a list of the cells that you want to be involved in

the function. (You can also do this using the mouse, which I'll explain later, but typing the function operator is easier for the purposes of this example.) Figure 17-2 shows the sample spreadsheet at this point.

FIGURE 17-2

Beginning to define a formula



NOTE **As you enter this value in the cell, you will see the function that you are entering mirrored in the Formula bar, which is a toolbar above the spreadsheet that simplifies editing existing formulas.**

Next, identify the values that you want to summarize, or perform an action on. This is most easily done using the mouse. To select the existing numeric values in a row, click the cell containing the cost value in the second column and hold down the left mouse button while dragging the cursor to the right to encompass the other two numeric values in that row. This is known as selecting a *range* of cells. As you select a range of cells, a dotted rectangle outlines the cells that you've selected, and the identifiers for those cells are listed following the opening parenthesis in the formula, as shown in Figure 17-3. (This is often referred to as a *selection box*.) As you can see in this figure, the cells in a spreadsheet are identified by their column and row — for example, B2 refers to the cell in column B, row 2. When represented in a formula, a range of cells is identified by referring to the first cell and the last cell in the selected range, separating their names by a colon. In this example, B2:D2 means "all cells between cell B2 and D2, including those two cells."

After selecting the cells for the function to act on, you can terminate the function by entering a closing parenthesis or simply press the Return/Enter key on your keyboard. Your spreadsheet calculates the value of the function that you've entered, displays that value, and moves the cursor to the next available cell, as shown in Figure 17-4. That was easy!

FIGURE 17-3

Selecting a range of cells for a formula to act on

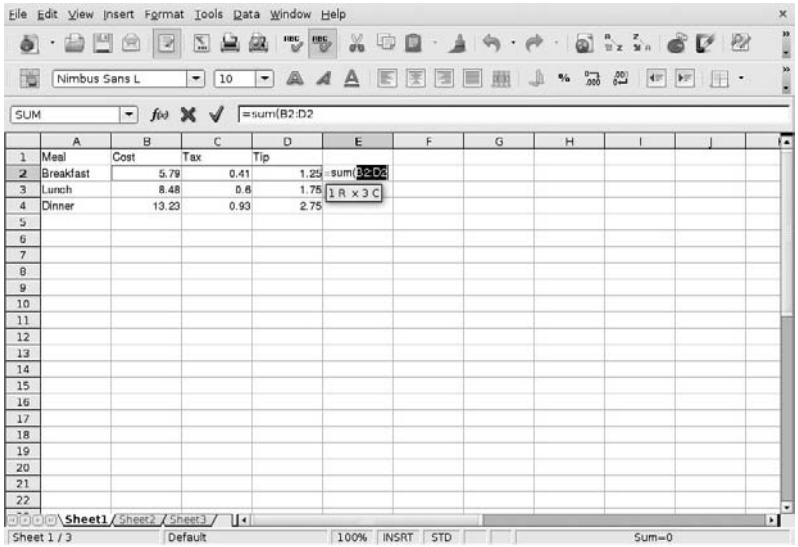
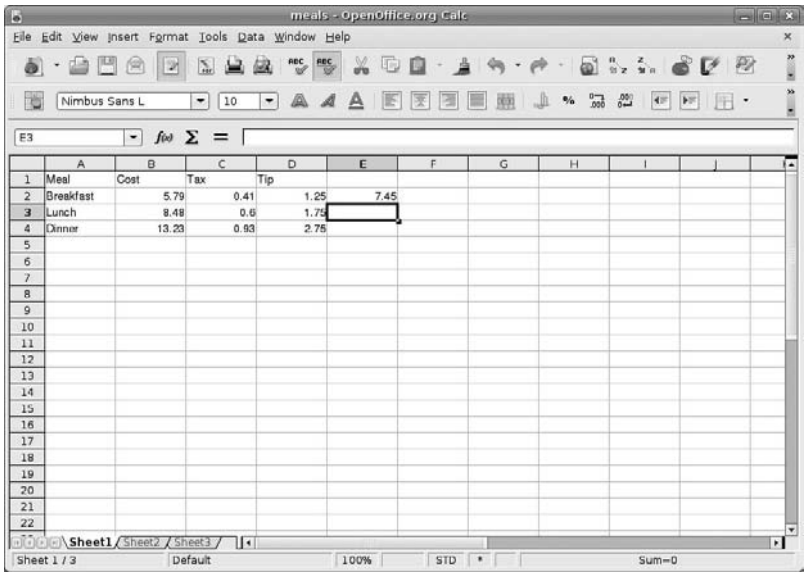


FIGURE 17-4

Displaying a calculated value



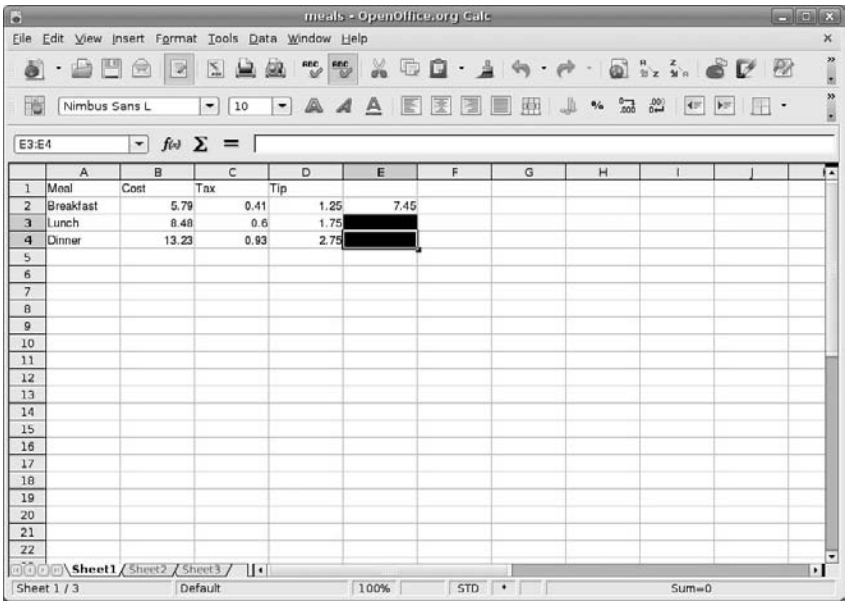
You could do the same thing in the next two rows, but there are easier ways to accomplish this that will help show some other basic spreadsheet capabilities. The easiest way to clone a calculated value to other cells in a column is to copy the contents of the first cell that contains the function and paste it into the other cells in that column. When you copy a calculated cell and paste it into multiple other cells at the same time, the spreadsheet figures out that you actually want to paste in a version of that function that applies to the rows in each column that you're pasting the function into. And people say that computer programs can't be smart!

To copy the function from the cell where you defined it, click that cell and select the Edit menu's Copy command. In some spreadsheets, a dashed line appears around any selected cell(s) to show what's selected, but that's not the case in OpenOffice.org Calc.

Next, select the range of cells that you want to paste that value into, in the same way that you selected a range of cells when defining your function. To do this, click the cell below the one that currently contains the formula and hold down the left mouse button while dragging the cursor down one row to encompass the same cell in the Dinner row, as shown in Figure 17-5.

FIGURE 17-5

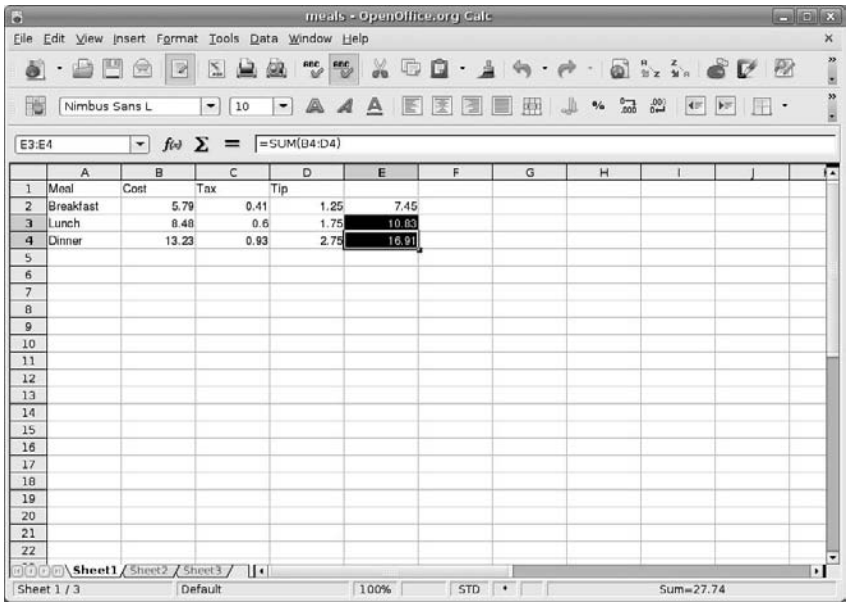
Selecting the range of cells to paste the function into



Finally, select the Edit menu's Paste command to paste the function into the selected range of cells. Voilà! The function is pasted into each cell and automatically calculates the right value for each row, as shown in Figure 17-6.

FIGURE 17-6

The updated spreadsheet with per-row calculations



I love it when a plan comes together! The only numeric calculation that remains is to create another cell that summarizes the calculated values in the per-meal summary column. Adding a cell for a summary is easy enough to do, in exactly the same way as you did before, which I'll summarize:

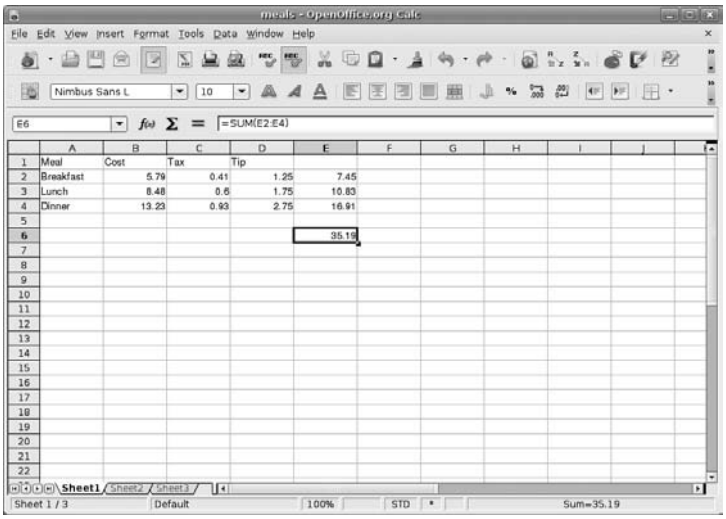
1. Select the cell in which you want to invoke a function.
2. Type `=sum(` to start defining the function.
3. Click the left mouse button, hold it down, and drag the selection box to select the three cells in the column that you're summarizing.
4. Enter the closing parenthesis to conclude the function and press Return to enter it and perform the calculation.

You can use any cell that you want. In my example, I skipped a row to visually make the summary cell stand out — your aesthetics and mine may differ. If all went according to plan, your spreadsheet should now look like the one shown in Figure 17-7. Congratulations — you've learned the basics of working with a spreadsheet!

Aside from the fact that you didn't have to do any of the math in this example (hooray!), one of the coolest things about a spreadsheet containing calculated values is that you can update any of the other cells that are involved in those calculations, and the spreadsheet will automatically update itself to reflect the new value in any cells that are based on the contents of the modified cell. Try it yourself — as all good software should do, *it just works*!

FIGURE 17-7

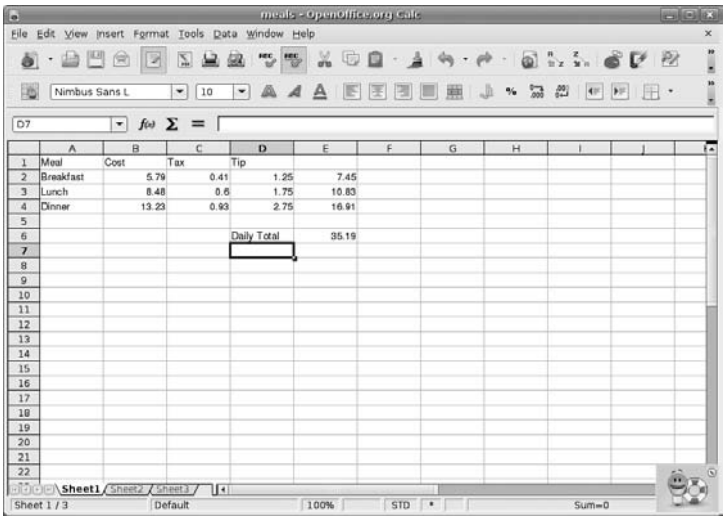
The sum of your previously calculated sums



Once the basic calculations are complete, you can add some labels to clarify the meaning of the values in the different fields. You already have row and column headings, but nothing explains the cell that displays the daily total. To enter text in a cell of the spreadsheet, click that cell to select it and type whatever you want. Figure 17-8 shows the sample spreadsheet after I've entered the value "Daily Total" in cell D6.

FIGURE 17-8

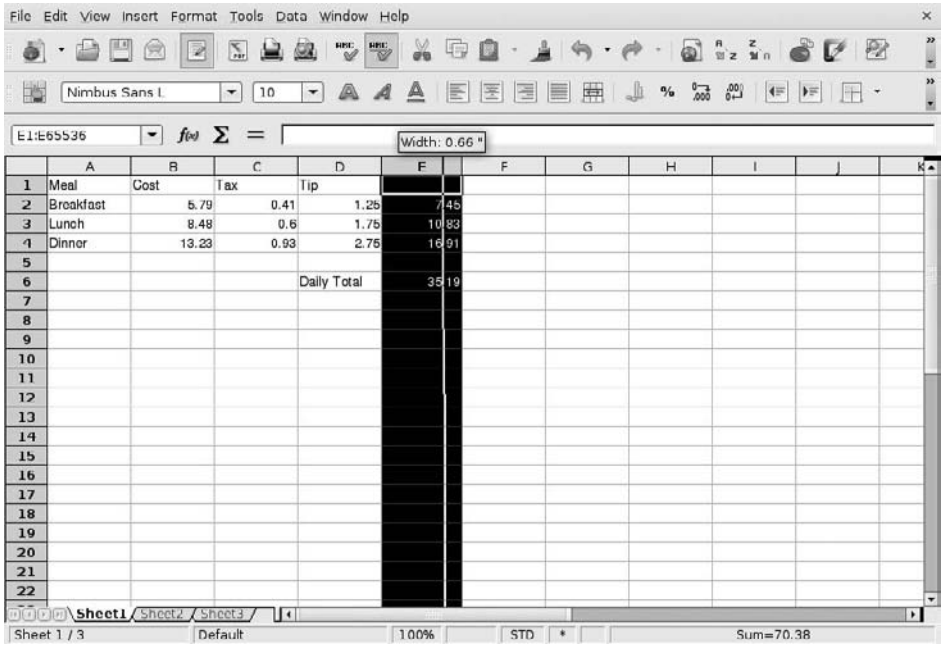
Adding a new text field to the spreadsheet



Some of the columns in the spreadsheet are a bit wider than they need to be. To fix this, you can adjust the column width. To adjust the width of a column in a spreadsheet, left-click the column heading (let's use column E as an example) to select it. Next, position the cursor over the line at the right edge of that column, hold down the left mouse button to select that edge, and drag it to the right to increase its size. As you drag the column edge, a tooltip appears that displays the current column width, as shown in Figure 17-9.

FIGURE 17-9

Resizing a column in the spreadsheet



Once you've resized the column so that all of your text is visible, release the left mouse button to complete the resize operation. The column displays on the screen with its new width, as shown in Figure 17-10.

Next, let's clean up the formatting a bit. It's normal to want to visually differentiate labels from data values, so you can make all of the headings and labels bold. You can do this easily by multi-selecting all of them and applying bold formatting. To do this, click on the row heading for row 1, hold down the Ctrl key to indicate that you want to select multiple items, click the column heading for column A, and also select the cell containing the "Daily Total" label. Your screen should look like Figure 17-11.

Next, click the Bold button in the formatting toolbar (the fat "A" icon). After you've applied the bold, it may be easier to see the formatting if you select (click on) another cell first, so that the bold cells are no longer selected. All selected items are formatted in bold, and your screen should look like the spreadsheet shown in Figure 17-12.

FIGURE 17-10

The resized column

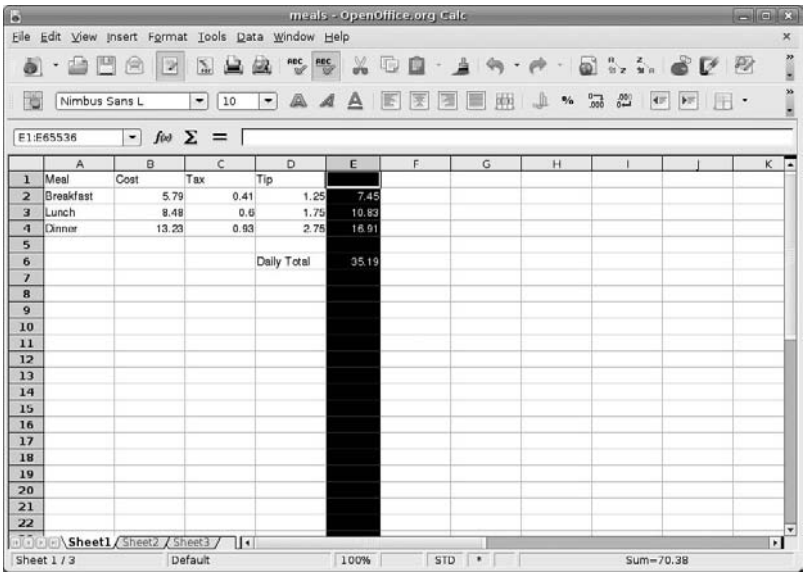


FIGURE 17-11

Multi-selecting all labels in the spreadsheet

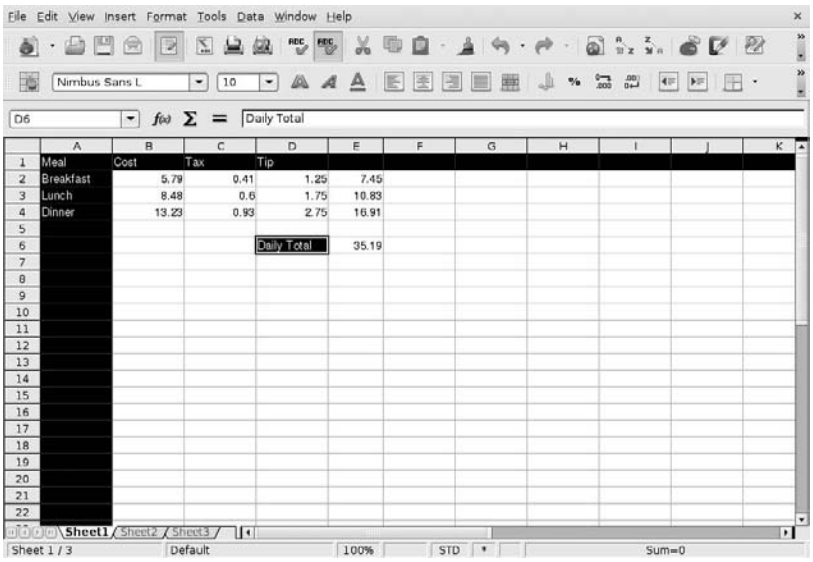
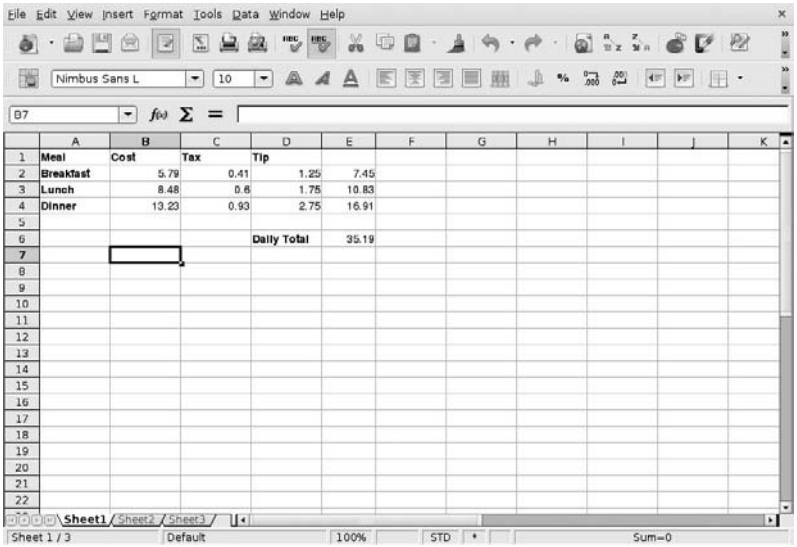


FIGURE 17-12

Updated formatting in the spreadsheet



The last thing to do after you've done any work is to save your file, which you can do by selecting the File ➪ Save command. Because your spreadsheet has only a temporary name at this point, whatever spreadsheet you're using will display a dialog that enables you to specify a name for your spreadsheet (and optionally navigate to a specific directory location in which to save the spreadsheet).

Of course, there's much more to today's spreadsheets than the simple actions described in this tutorial, but this should be enough to get you started. All of the spreadsheets discussed in this chapter provide excellent online help, and plenty of books are available that explain how to use spreadsheets. Although few books are available on Linux-specific spreadsheets (yet), OpenOffice.org Calc, KSpread, and Gnumeric are designed to be compatible with and similar to more common spreadsheets such as Microsoft Excel, so you should be able to get a lot of information out of almost any book on common spreadsheets.

Creating and Using OpenOffice.org Spreadsheets

As mentioned in the introduction to this chapter, spreadsheets are often as important to home computer users as they are to business users. This is especially true once a year for those of us in the United States, when tax time for anyone who itemizes deductions means trotting out all of your receipts for the past year, entering them into a spreadsheet, and totaling up anything that is legally deductible.

A recent poll by one of the Linux publications queried readers about what software they'd most like to see ported to Linux, and my vote went to TurboTax, because otherwise I'll have to have at least one Mac OS X machine around the house for the next 20 years or so. Until the TurboTax folks heed my plea (or, more

likely, recognize Linux as a growing market opportunity), spreadsheets will have to do for tracking my personal business expenses. Fortunately, Linux provides some excellent spreadsheet packages.

The spreadsheet component of the OpenOffice.org package is called OpenOffice.org Calc, although you start it by selecting the OpenOffice.org Spreadsheet menu item. The OpenOffice.org Office suite is a set of applications that were designed together, share extensive amounts of code, have a similar look and feel, and invoke each other to satisfy certain requirements. For example, as I discuss later in this chapter, bulk-importing existing spreadsheets into Calc using its Document Converter Wizard produces a report of the results of the import process. This file is created in the OpenOffice.org document format used by OpenOffice.org Writer, and can be displayed once the conversion process completes so that you can check for any errors or problems. You can't even select the Document Converter's Show log file button if OpenOffice.org Writer is not installed on your system. Fortunately, all of the OpenOffice.org Office suite is pre-installed on Ubuntu and Kubuntu desktop systems nowadays.

NOTE **I'll refer to the OpenOffice.org Spreadsheet package as Calc throughout the rest of this chapter because, well, that's its name. There is something to be said for using menu item labels such as Spreadsheets, Presentation, Database, and Word Processor to make it easier for new users to select the application associated with the type of document that they want to create, but I usually refer to programs by their actual names.**

As companies such as Microsoft have discovered, there's a lot to be said for a set of consistent, related applications that know how to leverage each other's capabilities to perform many of the data management tasks required in a modern, computer-oriented office. Unlike Microsoft Office, the OpenOffice.org suite is a high-quality, no-cost solution to your office software requirements. Although I've been a Gnumeric user for years (see the section on Gnumeric, later in this chapter), I tend to use all of the OpenOffice.org applications on my desktop systems nowadays. Because Gnumeric is much smaller and lighter-weight than OpenOffice.org, I still use Gnumeric on my laptops, which has come in very handy both times that I've needed a spreadsheet in a hotel room in recent years.

Installing Files for OpenOffice.org Calc

The OpenOffice.org software, including Calc, is installed by default if you install Ubuntu or Kubuntu from their Live CDs, from the Ubuntu Alternate Install CD using the Text Mode installer, and from the Ubuntu Alternate installer using the OEM Mode installer. If you want to add Calc or the entire OpenOffice.org suite to an Ubuntu server system that has a graphical user interface, you can install them using apt-get, aptitude, or Synaptic, as explained in Chapter 23, "Adding, Removing, and Updating Software."

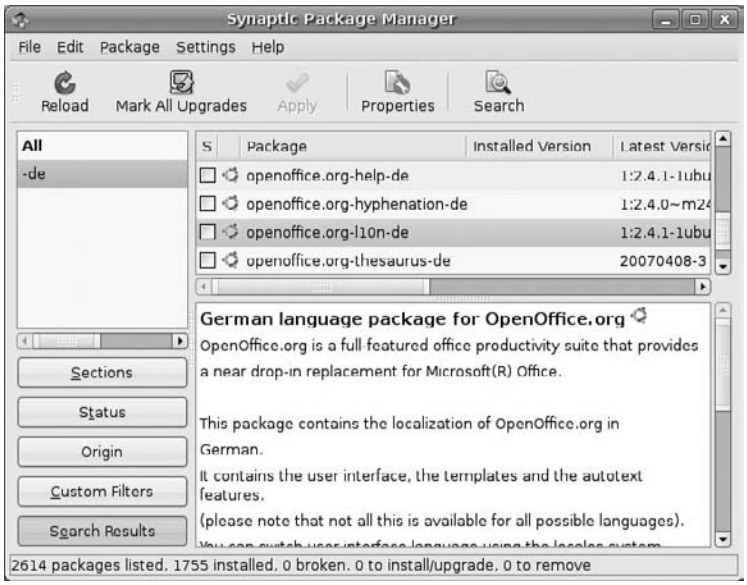
One of the things that differentiates Ubuntu Linux from other Linux distributions is its extreme dedication to the user, especially in terms of internationalization. If you are installing Calc for use when creating spreadsheets in languages other than English, you may also want to search for and install the appropriate hyphenation, localization, and thesaurus packages for whatever other locales you are creating spreadsheets for. This is especially important for working with spreadsheets that use non-U.S. currencies.

TIP **If you've already installed the localization and internationalization packages for any component of the OpenOffice.org suite, there is no need to install them again. Installing them once installs them as a shared resource that can be used by every component of the suite.**

You can easily find these localization and internationalization packages by looking for a string that represents the country code that you are looking for. For example, Figure 17-13 shows the Synaptic Package Manager, displaying specific German language support packages for OpenOffice.org that have been identified by searching for the string "-de."

FIGURE 17-13

Related language packages for OpenOffice.org



TIP **You can install complete language support for all applications (without translating the GNOME user interface itself) by installing the appropriate language support meta-package for your target locale(s). These packages have names of the form language-support-country-code (for example, language-support-de), and include the dictionaries for the specified locale. To fully translate the GNOME interface and all internationalized applications, select the appropriate language pack. The names of the language packs are of the form language-pack-country-code. Once you've installed multiple dictionaries, and so on, you can set the default language for your document using the Tools ➪ Options dialog's Language Settings ➪ Language drop-down list.**

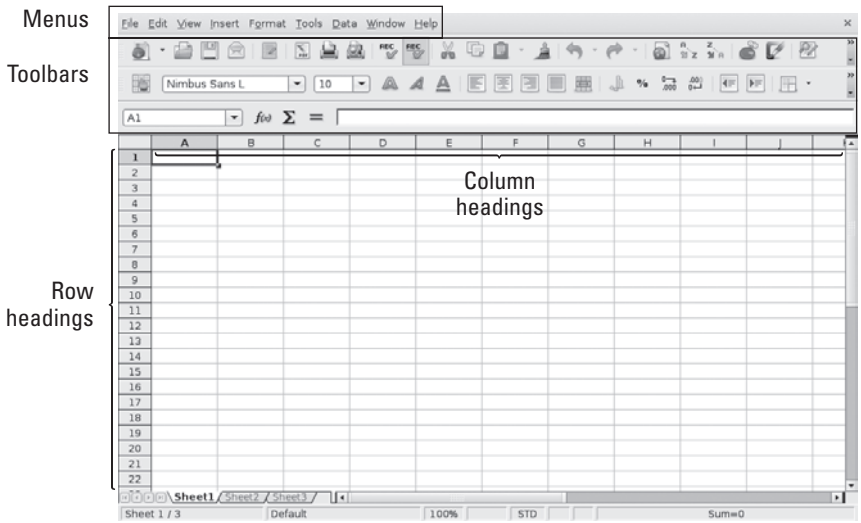
See Chapter 23, "Adding, Removing, and Updating Software," for more detailed information on installing applications using the Synaptic Package Manager on Ubuntu or the Adept tools on a Kubuntu system.

Starting Calc

You can start Calc from any command line by executing the command oocalc, but it is usually easier to start it from your GNOME menus. A menu item for Calc is provided in the Applications ➪ Office menu. To start Calc, select the OpenOffice.org Spreadsheet menu item. The application starts, displaying an empty document, as shown in Figure 17-14.

FIGURE 17-14

Starting Calc



Taking a Quick Tour of Calc

OpenOffice.org Calc provides excellent online help, as do all of the tools in the OpenOffice.org suite, so I won't bore you by walking you through each menu and entry that it contains. However, as a quick overview of the Calc interface, the callouts in Figures 17-14 and 17-15 highlight the following areas:

- **Menus:** The drop-down menus in Calc are organized much like the menus in Microsoft Excel to provide a familiar user experience. Menu items with an arrow at their right edge lead to other, related submenus.
- **Toolbars:** By default, Calc displays the Standard and Formatting toolbars, which provide one-click access to specific commands when you click on the appropriate icon, and a Formula Bar, in which you can view and edit any functions or formulas for selected cells. You can modify the list of icons displayed on a toolbar by clicking on the control for that toolbar, located at the far right. You can display other toolbars or deactivate the default toolbars by selecting the names of those toolbars from the View ➪ Toolbars submenu.
- **Row and column headings:** These headings identify the current row and column for the cells that you're viewing. They can be useful when manually entering cells or ranges in a formula.
- **Context-sensitive menu:** Right-clicking on any cell in a spreadsheet that you are working on in Calc displays the context-sensitive menu shown in Figure 17-15. This menu provides quick access to several cell manipulation commands.

Calc provides extensive online help, available by selecting OpenOffice.org Help from the Help menu, or by pressing the F1 key on your keyboard at any time. Figure 17-16 shows the main page for Calc's help, which provides both a searchable index and a list of popular, general topics from its initial display window (at right).

FIGURE 17-15

Calc's context-sensitive menu

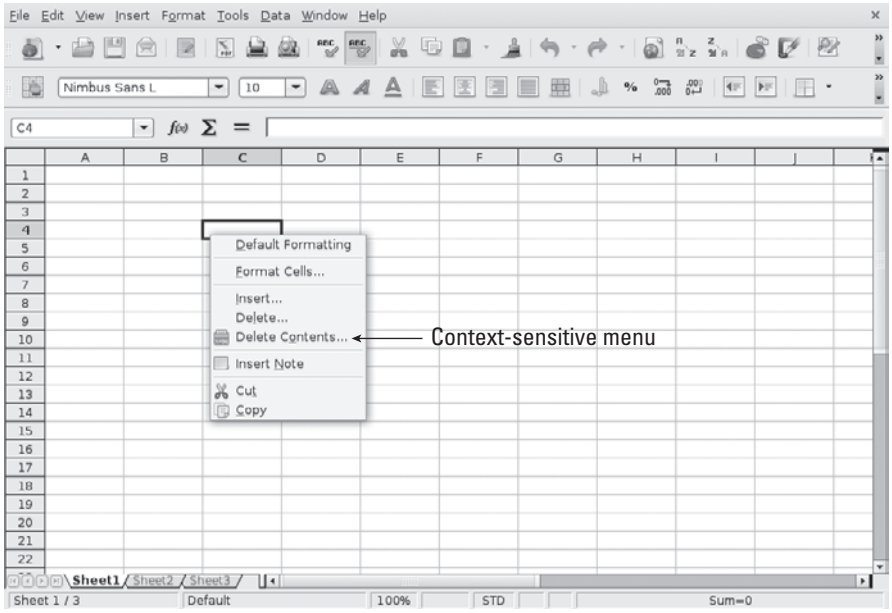
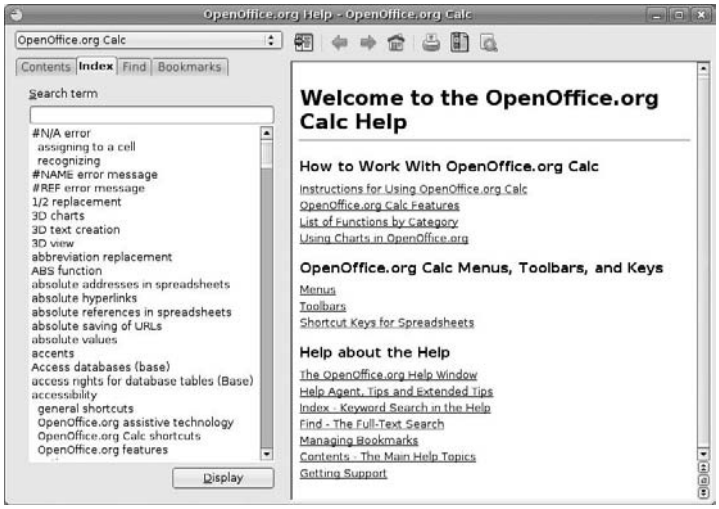


FIGURE 17-16

Calc's online help



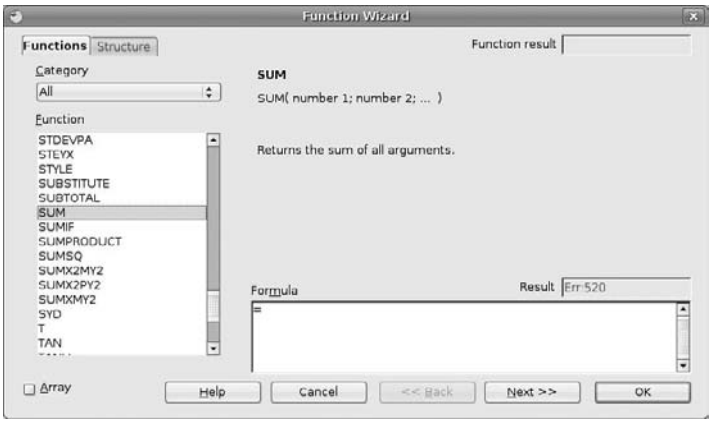
TIP **If you are having problems viewing the text for Calc's online help and OpenOffice.org Writer is not installed on your system for some reason, try installing Writer. This will force an update to the components that control online Help in the OpenOffice.org suite, and usually resolves the problem.**

Invoking Functions in Calc

The tutorial given earlier in this chapter explained how to manually enter a function, which is fine if you are using a simple function that you happen to know the name of. Calc supports an extensive collection of predefined functions that you can easily select from the Insert ➪ Function dialog using the mouse. Figure 17-17 shows this dialog with the SUM function selected from the complete list of available formulas.

FIGURE 17-17

Available functions in Calc



As you can see, selecting any function from the Function dialog displays abbreviated help for that function in the pane at the right of the dialog — in this case, "Returns the sum of all arguments."

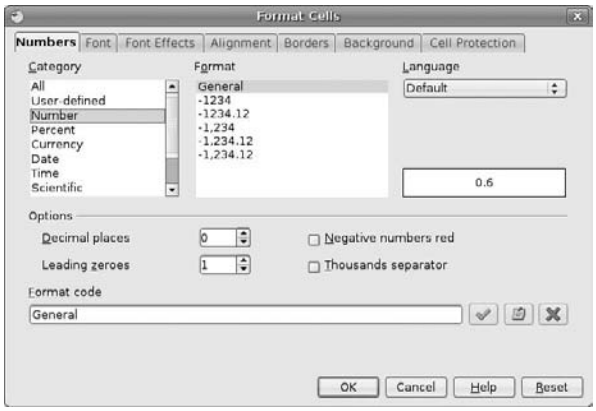
Specifying the Type of Data in a Cell

The spreadsheet tutorial in this section used fairly straightforward types of data that Calc could easily recognize — sequences of characters from the alphabet and spaces are probably words, whereas numbers and decimal points are probably numeric values. Knowing the type of data that is contained in a cell (generally referred to as its *data type* or *data format*) is fairly important, especially when you subsequently want to invoke functions that calculate values from the contents of various cells.

When you enter values in a spreadsheet without specifying their type, spreadsheets such as Calc assign them a general data format, which performs a best-guess of their data format based on the type of characters that they contain. However, you can also manually specify the format of any cell (or cells) by selecting those cells and using the commands on the Format ➪ Cells dialog, shown in Figure 17-18.

FIGURE 17-18

Specifying cell data formats in Calc



As an example, note that cell C3 in your example spreadsheet (most recently shown in Figure 17-12) doesn't display the traditional number of decimal places that you'd expect to see in a monetary value. This is because Calc's general data format recognized it as a decimal number, but didn't have any way of knowing that this was a monetary value. To correct this, select that cell, select the Format ➪ Cells command, and click Currency in the pane at left. This displays a list of possible representations for that data format, as shown in Figure 17-19.

Note that two decimal places is the default for this data type, which is what you are trying to correct. Because this would do the trick, click OK to set cell C3 as a currency value and close the Format Cells dialog. This updates the spreadsheet, which now shows the correct number of decimal places for cell C3, as shown in Figure 17-20.

FIGURE 17-19

Specifying that a cell contains a currency value

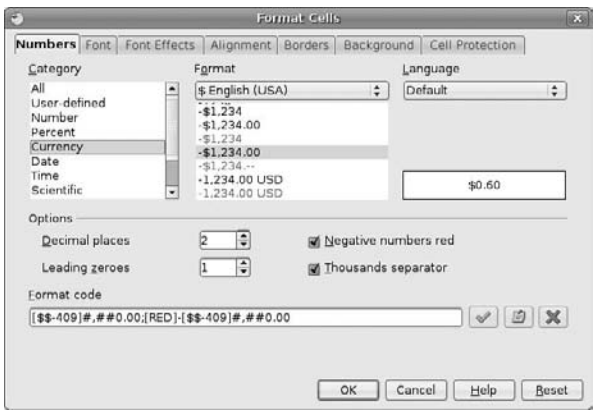


FIGURE 17-20

The updated spreadsheet showing the contents of cell C3 as a currency value (compare with Figure 17-12)

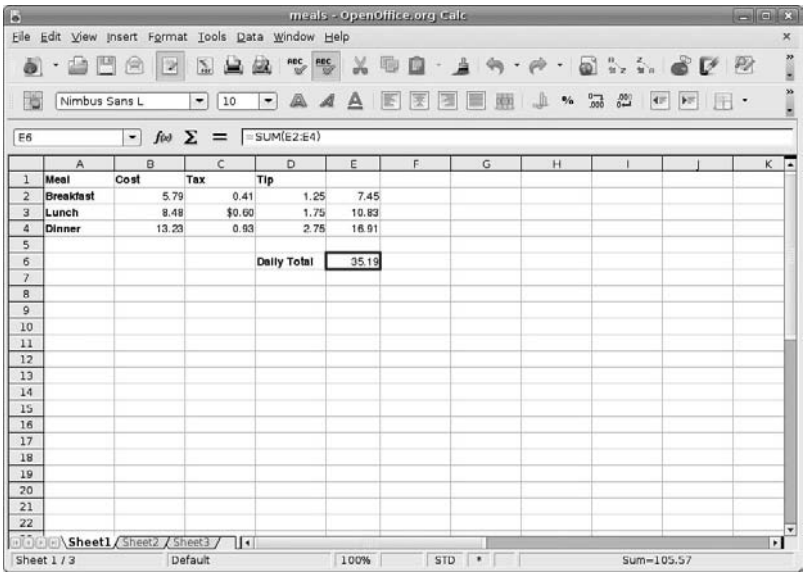
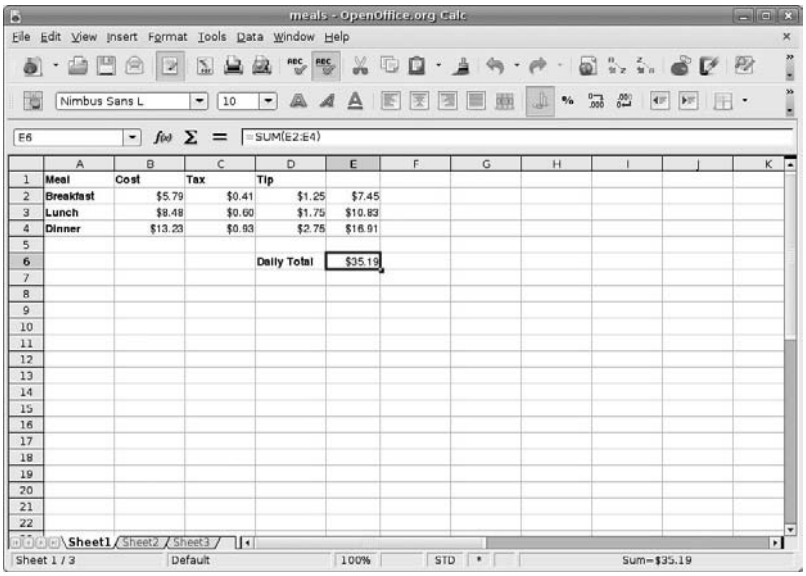


FIGURE 17-21

The updated spreadsheet showing uniform currency values



This is fine, except that Calc now recognizes that cell as a currency value but all other numeric cells are still general numeric values, which makes the spreadsheet look a bit odd. You can correct this by multi-selecting all of the numeric values in the spreadsheet (hold down the Ctrl key, click and hold the left mouse button on cell B2, drag the selection box to encompass all cells from B2 through E4, and then also click cell E6 before releasing the Ctrl key), displaying the Format Cells dialog again, selecting Currency, and clicking OK. Voilà! The cells in the spreadsheet are now consistent, as shown in Figure 17-21.

Importing Existing Spreadsheets into Calc

Linux folks like me tend to be Linux-centric, but that doesn't mean that we can completely ignore reality. When you're working with existing spreadsheets (and any other office-related documents, to be honest), the chances are pretty good that you created them using more common software packages, such as Microsoft Excel, that run on other operating systems, such as Mac OS X and Microsoft Windows. The folks who wrote and support the OpenOffice.org suite and OpenOffice.org Calc couldn't ignore this reality either, so Calc provides built-in support for opening Microsoft Excel spreadsheets and workbooks.

NOTE **If you have a lot of existing spreadsheets and are wondering how to transfer them to your new Ubuntu system, see Bonus Chapter 2, "File Transfer and Sharing on Ubuntu," for detailed information on doing just that. Don't worry — it's really quite easy to do using a network connection or even a traditional mechanism such as a floppy (if you still have a floppy drive) or a flash drive, or burning a CD.**

For example, to open an existing Excel spreadsheet, simply select the File ➪ Open command. By default, the file selection dialog that opens enables you to select files with any extension, including the .xls extension used by Microsoft Excel spreadsheets.

Supported Import Formats

The same mechanism described in the previous section can be used to open files in any other format supported by Calc. Calc can also open spreadsheets or other data files that have been saved in any of the following formats:

- **Comma-Separated Value** (`.csv`)**:** A standard text-based interchange format, where each row appears on a single line and each field value is separated from the next by a comma. This exchange format does not preserve the calculations required to produce cell values, only the values of the cells themselves.
- **Data Interchange Format** (`.dif`)**:** Another standard text-based interchange format that does not preserve the calculations required to produce cell values, only the values of the cells themselves.
- **Lotus 1-2-3** (`.123`, `.wk1`, `.wks`)**:** The spreadsheet formats used by various versions of Lotus 1-2-3 spreadsheet.
- **Microsoft Excel** (`.xls`, `.xlsx`, `.xlt`, `.xlw`, `.xml`)**:** All of the common spreadsheet formats used by versions of Excel.
- **Multiplan** (`.sylk`)**:** The spreadsheet format used by the old Multiplan spreadsheet.
- **Open/StarOffice Calc** (`.ods`, `.sdc`)**:** The open document spreadsheet format used by Calc and other forward-thinking (i.e., open) spreadsheets.
- **Quattro Pro** (`.wb1`, `wb2`, `.wb3`)**:** The spreadsheet formats used by various versions of Borland's Quattro Pro spreadsheet.
- **Pocket Excel** (`.pxl`)**:** The spreadsheet format used by versions of Microsoft Excel for PDAs.
- **Xbase** (`.dbf`)**:** The database file format used by a variety of Dbase III–compatible applications such as Microsoft FoxPro, Dbase III and later, and so on.

As you can see, the open source developers who are responsible for Calc have made a very significant effort to ensure that you can import any of your existing spreadsheet data files into Calc, so that you're not losing anything other than your dependence on your previous operating system.

Importing Spreadsheets Using the Document Converter

As you can see from the previous section, Calc can open spreadsheets produced by many other spreadsheet applications. It can also open spreadsheet files with Microsoft Excel extensions directly. After working on them, you can then save them directly in the new, open OpenDocument spreadsheet format (`http://en.wikipedia.org/wiki/OpenDocument`) that is used by Calc.

TIP **If you're actively sharing spreadsheets with people who are still using commercial spreadsheet application such as Excel, you'll want to use the File menu's Save As command to specify that you want to save them in their original format, such as Excel.**

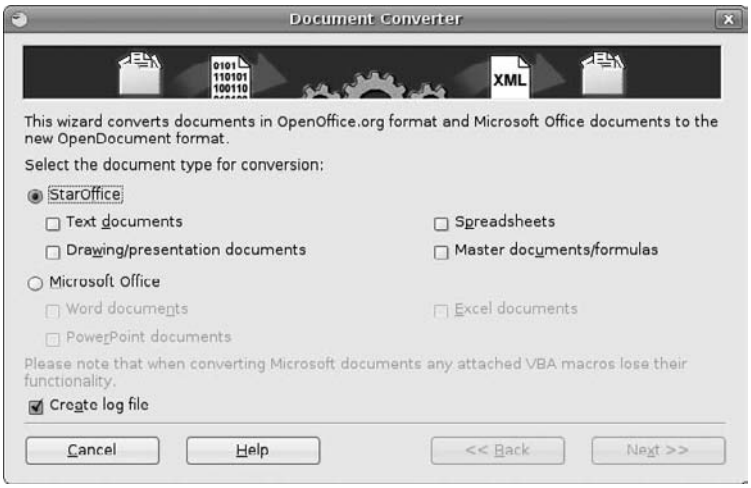
Repeating this process is fine if you have only one or two spreadsheets that you want to update, but if you are moving all of your work to Ubuntu, you may find that you have hundreds of spreadsheets that you want to update to Calc format. Fortunately, Calc provides a wizard for this that will not only import all of your Microsoft Excel spreadsheets, but also will import and convert any existing or attached templates that it detects.

To batch-import your existing Microsoft Excel (or older OpenOffice.org) documents, do the following:

1. Select File ➪ Wizards ➪ Document Converter. The dialog shown in Figure 17-22 displays.

FIGURE 17-22

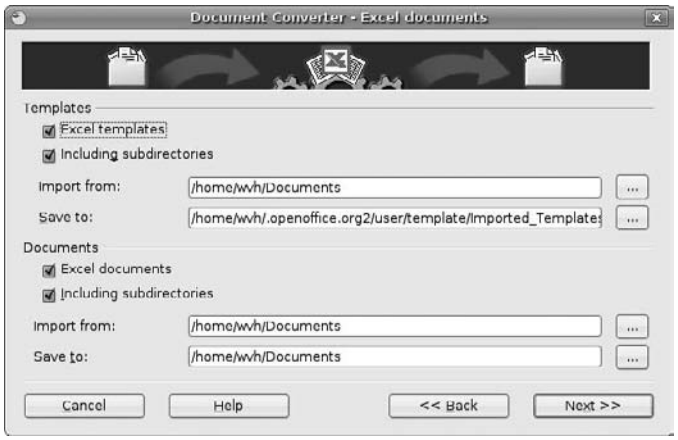
Specifying the input format when converting documents



2. To batch import Microsoft Excel files and templates, select the Microsoft Office radio button and the Excel Documents checkbox (grayed out in the figure, but if you select the Microsoft Office radio button, the Excel Documents checkbox will become available). Click Next to proceed. The dialog shown in Figure 17-23 displays.

FIGURE 17-23

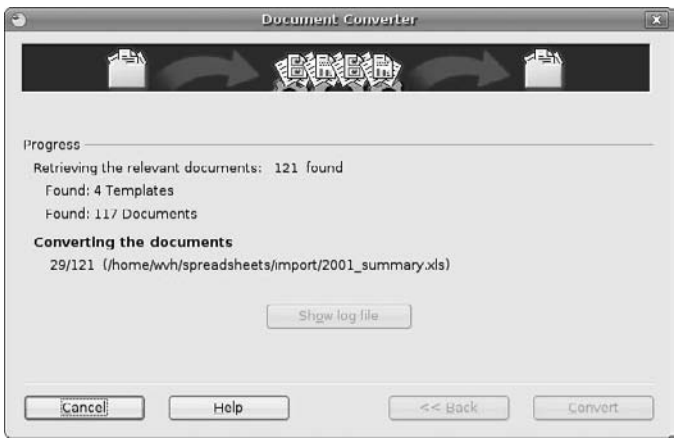
Specifying document and template locations when importing



3. Specify the directories in which you want the wizard to search for the specified document types, and the locations to which converted documents should be saved. You may want to update the location to which converted templates should be saved — I typically use my home directory, and subsequently move the files around manually. Once you've checked the directory locations, click Next to proceed. A status dialog displays, summarizing the conversions that are about to take place. Click Convert to begin the conversion process. The dialog shown in Figure 17-24 displays as the conversion process progresses. When the conversion process completes, the Close and Show log file buttons are activated on the dialog.

FIGURE 17-24

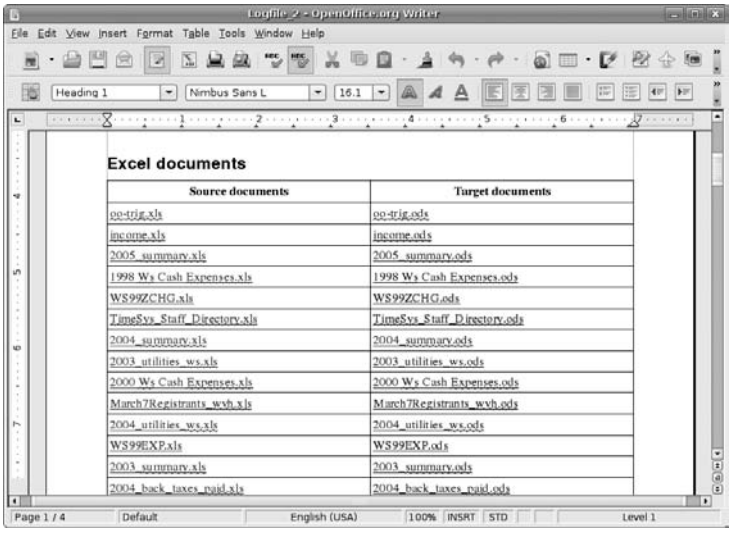
Monitoring the document conversion process



Congratulations — the conversion process is complete! Your converted spreadsheets and templates are ready for use in Calc. Before clicking Close, it's a good idea to click the Show log file button, which opens the log file in OpenOffice.org Writer. (This button will not be selectable if OpenOffice.org Writer is not installed on your system.) As shown in Figure 17-25, the log file for the conversion process lists successful conversions, and also identifies spreadsheets that should be fine-tuned or corrected before they can be used with Calc (if there are any). This is often the case with spreadsheet template files or spreadsheets containing complex macros, which may contain statements that are not completely compatible with anything other than Microsoft's own Visual Basic.

FIGURE 17-25

A log file of the spreadsheet conversion process



Creating and Using Gnumeric Spreadsheets

Gnumeric, a component of the GNOME Office suite, was one of the first spreadsheets available for Linux and, like a fine wine, has only continued to improve with age. The GNOME Office suite is a collection of office applications that depend on the GIMP Toolkit (GTK) graphics library for many of their graphical controls. These have been semantically collected into the GNOME Office suite through their allegiance to GTK and the need for the GNOME folks to have an office suite just like the KDE folks, not necessarily because of any internal similarities or common functionality beyond GTK itself. Although OpenOffice.org Calc is installed by default on Ubuntu desktop systems, you may find that you prefer using a lighter-weight spreadsheet such as Gnumeric, and this may be especially true on systems with limited disk space or a limited amount of memory. Linux is all about choice — and Gnumeric is a fine choice for a powerful and highly usable spreadsheet if there's something that you don't like about OpenOffice.org Calc. You can also install Gnumeric on Kubuntu systems if you want to, although KSpread is an equally powerful (and preinstalled) spreadsheet that is preferred by most KDE users.

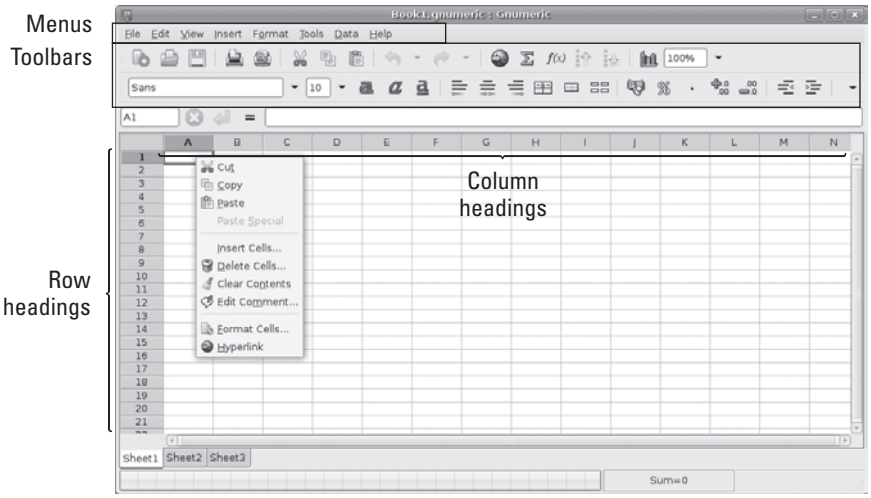
NOTE **Gnumeric was installed by default in earlier versions of the Ubuntu desktop distribution, but is not installed by default on Ubuntu 8.04. You can install Gnumeric on your system by using any of the software update utilities (`apt-get`, `aptitude`, or `Synaptic`) discussed in Chapter 23, "Adding, Removing, and Updating Software." The packages that contain Gnumeric are the gnumeric and gnumeric-common packages. If you are installing Gnumeric, you may also want to install the gnumeric-doc package, which provides the entire Gnumeric manual in GNOME help format, and the gnumeric-plugins-extra package, which provides some more advanced plug-ins for Gnumeric, including support for Python and Perl plug-ins.**

Starting Gnumeric

After you install Gnumeric on your system, a menu item for it is present on the Applications ➪ Office menu. To start Gnumeric, select this menu item. The application starts, displaying an empty document, as shown in Figure 17-26.

FIGURE 17-26

Starting Gnumeric



Taking a Quick Tour of Gnumeric

Gnumeric provides excellent online help, so I won't bore you by walking you through each menu and entry that it contains. However, as a quick overview of the Gnumeric interface, the callouts in Figure 17-26 highlight the following areas:

- **Menus:** The drop-down menus in Gnumeric are organized much like the menus in Microsoft Excel to provide a familiar user experience. Menu items with an arrow at their right edge lead to other, related submenus.
- **Toolbars:** By default, Gnumeric displays the Standard and Format toolbars, which provide one-click access to specific commands when you click on the appropriate icon. Depending on the size of the window in which Gnumeric is being displayed, you can access other commands that didn't

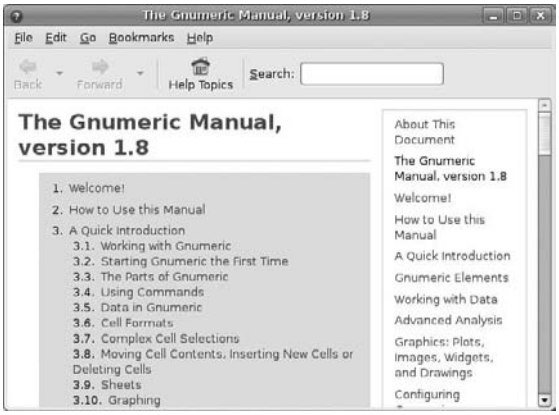
fit on the current view of the toolbar by clicking on the down arrow located at the far right. You can also display other toolbars or deactivate the default toolbars by selecting toolbars from the View ➪ Toolbars submenu.

- **Row and column headings:** These headings identify the current row and column for the cells that you're viewing. They can be useful when manually entering cells or ranges in a formula.
- **Context-sensitive menu:** Right-clicking on a cell in any spreadsheet that you are working on in Gnumeric displays the context-sensitive menu shown in Figure 17-26. This menu provides quick access to several cell manipulation commands.

Gnumeric provides extensive online help, available by selecting Contents from the Help menu, or by pressing the F1 key on your keyboard at any time. Figure 17-27 shows the main page for Gnumeric's online help, which provides both a searchable index and a list (at right) of popular, general topics from its initial display window.

FIGURE 17-27

Gnumeric's online help



You can also select the Help menu's Gnumeric on the Web menu item to display Gnumeric's home page on the Internet, as shown in Figure 17-28, from which you can always access a variety of online documentation and other resources for working with Gnumeric.

Invoking Functions in Gnumeric

The tutorial in the first section of this chapter explained how to manually enter a function, which is fine if you are using a simple function that you happen to know the name of. Gnumeric supports an extensive collection of functions that you can easily select from the Insert ➪ Function dialog using the mouse. The functions that you can select in this dialog are organized into multiple categories to make it easier to locate related functions whose exact names you may not know; you can also select the All Functions (long list) category to scroll through the entire list of functions that are available in Gnumeric. Figure 17-29 shows this dialog with the sum function selected from the Mathematics category.

As you can see from this figure, selecting any function from the Function dialog displays abbreviated help for that function in the pane at the bottom of the dialog (in this case, "SUM computes the sum of all the values . . .").

FIGURE 17-28

Gnumeric's home page on the Web



FIGURE 17-29

Available functions in Gnumeric



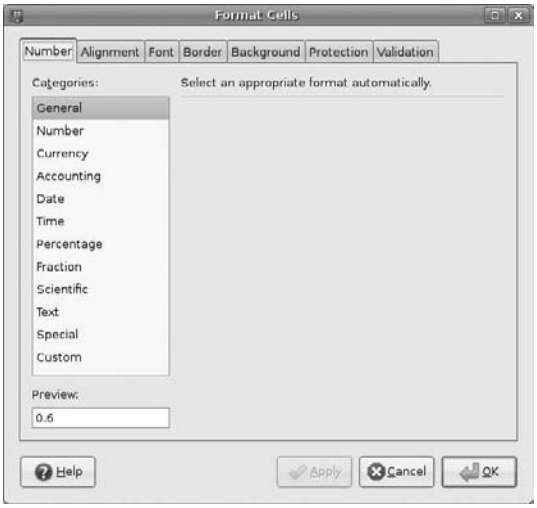
Specifying the Type of Data in a Cell

The spreadsheet tutorial in this section used fairly straightforward types of data that Gnumeric could easily recognize — sequences of characters from the alphabet and spaces are probably words, while numbers and decimal points are probably numeric values. Knowing the type of data that is contained in a cell (generally referred to as its *data type* or *data format*) is fairly important, especially when you subsequently want to invoke functions that calculate values from the contents of various cells.

When you enter values in a spreadsheet without specifying their type, spreadsheets such as Gnumeric assign them a general data format, which performs a best guess of their data format based on the type of characters that they contain. However, you can also manually specify the format of any cell (or cells) by selecting those cells and using the commands on the Format ➪ Cells dialog, shown in Figure 17-30.

FIGURE 17-30

Specifying cell data formats in Gnumeric



As an example, note that cell C3 in our example spreadsheet (most recently shown in Figure 17-12) doesn't display the traditional number of decimal places that you'd expect to see in a monetary value. This is because Gnumeric's general data format recognized it as a decimal number, but didn't have any way of knowing that this was a monetary value. To correct this, select that cell, select the Format ➪ Cells command, and click Currency in the pane at the left. This displays a list of possible representations for that data format, as shown in Figure 17-31.

Click the up arrow beside the Decimal Places label to specify two decimal places for this cell. You may also want to select the dollar sign as a symbol to precede the value (or a similar symbol, based on your geographic or national location) to visually identify this cell as a currency value rather than simply a decimal number. Once you've made your changes, click OK to set cell C3 as a currency value and close the Format Cells dialog. This updates the spreadsheet, which now shows the correct number of decimal places and a leading dollar sign for cell C3, as shown in Figure 17-32.

FIGURE 17-31

Specifying that a cell contains a currency value

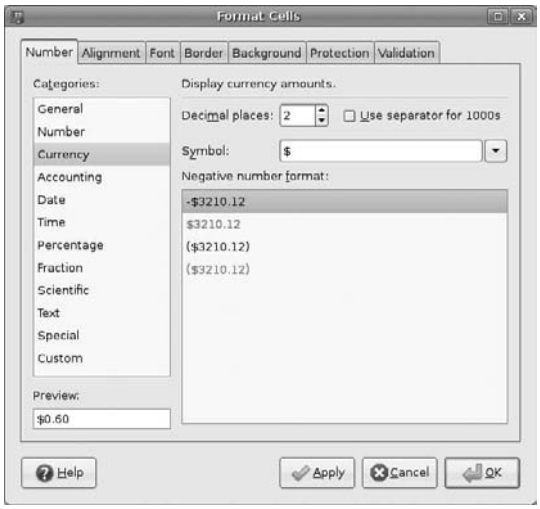
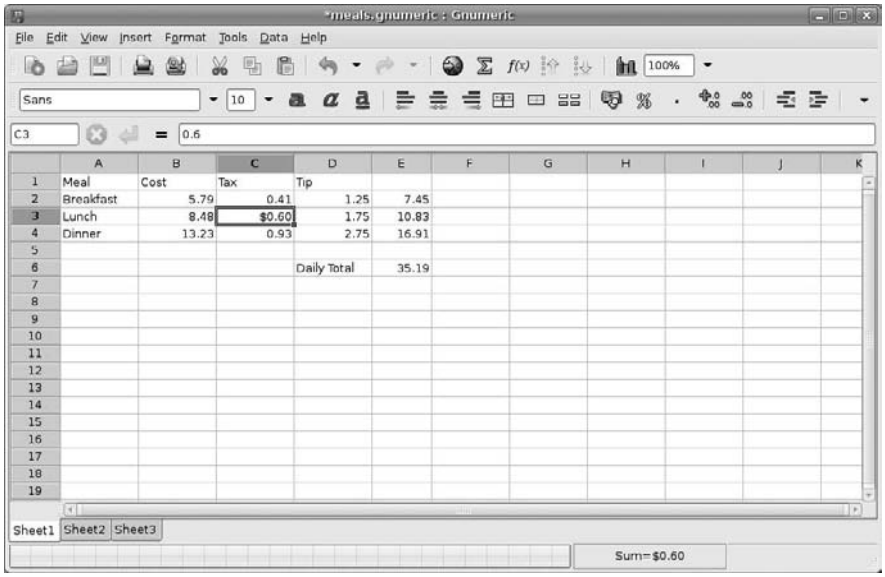


FIGURE 17-32

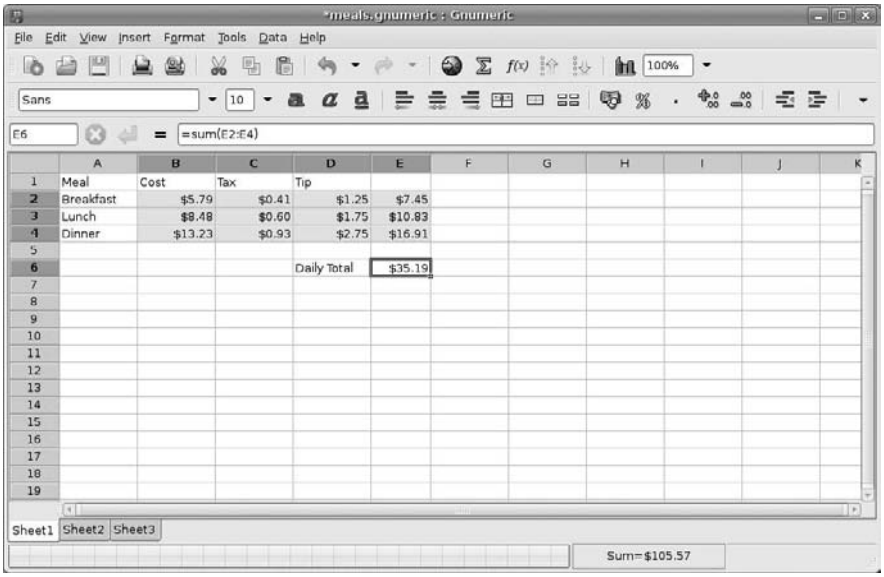
The updated spreadsheet showing C3 as a currency value



This is fine, except that Gnumeric now recognizes that cell as a currency value but all other numeric cells are still general numeric values, which makes the spreadsheet look a bit odd. You can correct this by multi-selecting all of the numeric values in the spreadsheet (hold down the Ctrl key, click and hold the left mouse button on cell B2, drag the selection box to encompass all cells from B2 through E4, and then also click cell E6 before releasing the Ctrl key), displaying the Format Cells dialog again, selecting Currency, making the same Decimal Place and Symbol settings, and clicking OK. Voilà! The spreadsheet is now consistent, as shown in Figure 17-33.

FIGURE 17-33

The updated spreadsheet showing currency values



Importing Existing Spreadsheets into Gnumeric

Linux folks like me tend to be Linux-centric, but that doesn't mean that we can completely ignore reality. When you're working with existing spreadsheets (and any other office-related documents, to be honest), the chances are pretty good that you created them using more common software packages, such as Microsoft Excel, that run on other operating systems, such as Mac OS X and Microsoft Windows. The folks who wrote and support Gnumeric couldn't ignore this reality either, so Gnumeric provides built-in support for opening Microsoft Excel spreadsheets and workbooks.

NOTE **If you have a lot of existing spreadsheets and are wondering how to transfer them to your new Ubuntu system, see Bonus Chapter 2, "File Transfer and Sharing on Ubuntu," for detailed information on doing just that. Don't worry — it's really quite easy to do using a network connection or even a traditional mechanism such as a floppy or a CD.**

For example, to open an existing Excel spreadsheet, simply select the File ➪ Open command. By default, the file selection dialog that opens enables you to select files with any extension, including the .xls extension used by Microsoft Excel spreadsheets.

The same mechanism is used to open files in any other format supported by Gnumeric. Gnumeric can open spreadsheets or other data files that have been saved in any of the following formats:

- **Applix** (`.as`)**:** The spreadsheet files produced by the Applix Office suite, which was discussed in Chapter 16.
- **Comma-Separated Value** (`.csv`)**:** A standard text-based interchange format, where each row appears on a single line and each field value is separated from the next by a comma. This exchange format does not preserve the calculations required to produce cell values, only the values of the cells themselves.
- **Data Interchange Format** (`.dif`)**:** Another standard text-based interchange format that does not preserve the calculations required to produce cell values, only the values of the cells themselves.
- **Gnumeric** (`.gnumeric`)**:** Gnumeric's own format.
- **HTML:** Tabular data in standard HyperText Markup Language format.
- **Linear and Integer Program Format** (`.mps`)**:** A text-based interchange format developed by IBM in the sixteenth century. This exchange format does not preserve the calculations required to produce cell values, only the values of the cells themselves, leaving the calculations to the monks who were illuminating the punch cards.
- **Lotus 1-2-3** (`.123`, `.wk1`, `.wks`)**:** The spreadsheet formats used by various versions of Lotus 1-2-3 spreadsheet.
- **Microsoft Excel** (`.xls`, `.xslx`)**:** The data format(s) used by Excel 2007 (xslx), and Excel 2003, and earlier Excel spreadsheets (xsl).
- **Multiplan** (`.sylk`)**:** The data format used by the old Multiplan spreadsheet.
- **Oleo** (`.mps`)**:** The data format used by the GNU project's Oleo spreadsheet (`www.gnu.org/software/oleo/oleo.html`), which was last updated/released in 2000.
- **OpenOffice.org/StarOffice Calc** (`.ods`, `.sxc`)**:** ODS is the open document spreadsheet format used by OpenOffice.org Calc and other forward-thinking (i.e., open) spreadsheets. SXC is the spreadsheet format used by StarOffice.
- **Quattro Pro** (`.wb1`, `wb2`, `.wb3`)**:** The spreadsheet formats used by various versions of Borland's Quattro Pro spreadsheet.
- **Plan Perfect** (`.pln`)**:** The data format used by the PlanPerfect component of Corel's WordPerfect Office suite.
- **Text** (`.txt`)**:** Text files that can be parsed in some way to convert them to the row-and-column model used by spreadsheets.
- **Xbase** (`.dbf`)**:** The database file format used by a variety of Dbase III–compatible applications such as Microsoft FoxPro, Dbase III and later, and so on.

As you can see, the open source developers who are responsible for Gnumeric have made a very significant effort to ensure that you can import any of your existing spreadsheet data files into Gnumeric, so that you're not losing anything other than your dependence on your previous operating system.

Creating and Using KSpread Spreadsheets

KSpread (`www.koffice.org/kspread/`) is the spreadsheet component of the KDE KOffice office suite (`www.koffice.org`). KSpread is installed by default on Kubuntu systems running KDE 3.5.*x*, but is not installed by default on Kubuntu systems running KDE 4.*x* (at least not at the time of this writing). KSpread provides a lighter-weight alternative to OpenOffice.org Calc as well as tight integration with the other components of the KOffice suite. This integration makes it easy to embed KSpread spreadsheets or charts in KWord documents, KPresenter presentations, and so on. KSpread has been actively under development since the late 1990s, and is a nice, full-featured spreadsheet package.

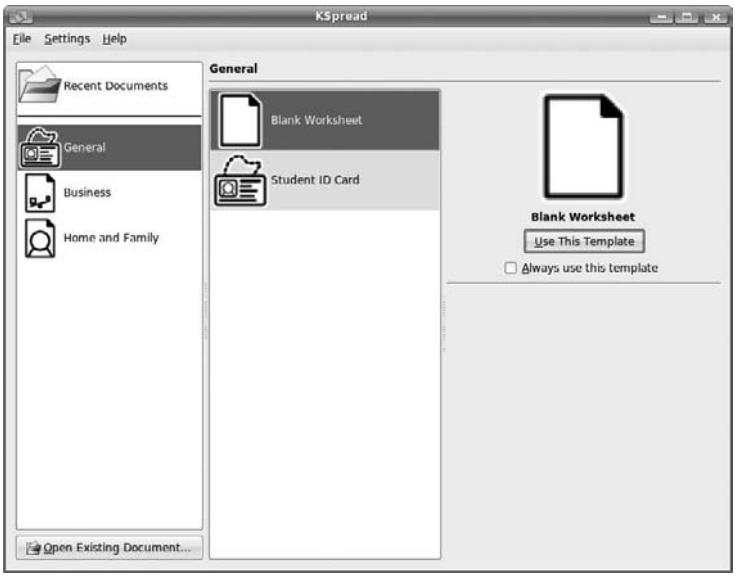
NOTE **KSpread can also be installed on GNOME-based Ubuntu desktop systems, but I think that the number of people who would want to do that is relatively small. Installing KSpread on a GNOME-based Ubuntu system means that you'd also have to install all of the underlying KDE components that it uses. I've used KSpread on and off on KDE-based systems for years, but using it on a GNOME-based system is not for me, though your mileage may vary. For information about installed KSpread on Ubuntu Desktop systems, see Chapter 23, "Adding, Removing, and Updating Software."**

Starting KSpread

To start KSpread on a Kubuntu system using KDE 3.5, select the K Menu ➪ Office ➪ KSpread spreadsheets menu item. (If you're using KSpread on an Ubuntu system for some reason, select the Applications ➪ Office ➪ KSpread menu item.) When first started, KSpread displays its spreadsheet wizard, as shown in Figure 17-34.

FIGURE 17-34

KSpread's spreadsheet wizard



KSpread's spreadsheet wizard enables you to select from different templates for any new spreadsheet that you are creating, and also provides the Recent Documents button and Open Existing Spreadsheet dialog to make it easy to reopen a KSpread document that you've worked on before.

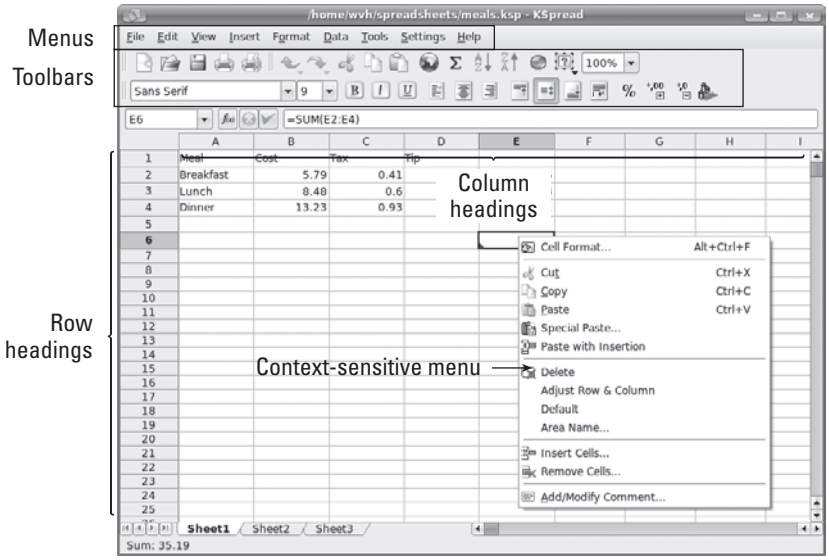
Taking a Quick Tour of KSpread

Like most KDE applications and certainly all of the KOffice suite, KSpread provides excellent online help, so I won't bore you by walking you through each menu and entry that it contains. However, as a quick overview of the KSpread interface, the callouts in Figure 17-35 highlight the following areas:

- **Menus:** The drop-down menus in KSpread are organized much like the menus in Microsoft Excel to provide a familiar user experience. Menu items with an arrow at their right edge lead to other, related submenus.
- **Toolbars:** By default, KSpread displays its File, Edit, and Format toolbars, which provide one-click access to specific commands when you click the appropriate icon. Depending on the size of the window in which KSpread is being displayed, you can access other commands that didn't fit on the current view of the toolbar by clicking the double right arrow located at the far right. You can also display other toolbars or deactivate the default toolbars by selecting those toolbars from the Settings ➪ Toolbars submenu.
- **Row and column headings:** These headings identify the current row and column for the cells that you're viewing. They can be useful when you are manually entering cells or ranges in a formula.
- **Context-sensitive menu:** Right-clicking on a cell in any spreadsheet that you are working on in KSpread displays the context-sensitive menu that is also shown in Figure 17-35. This menu provides quick access to several cell, row, and column manipulation commands.

FIGURE 17-35

Portions of the KSpread screen

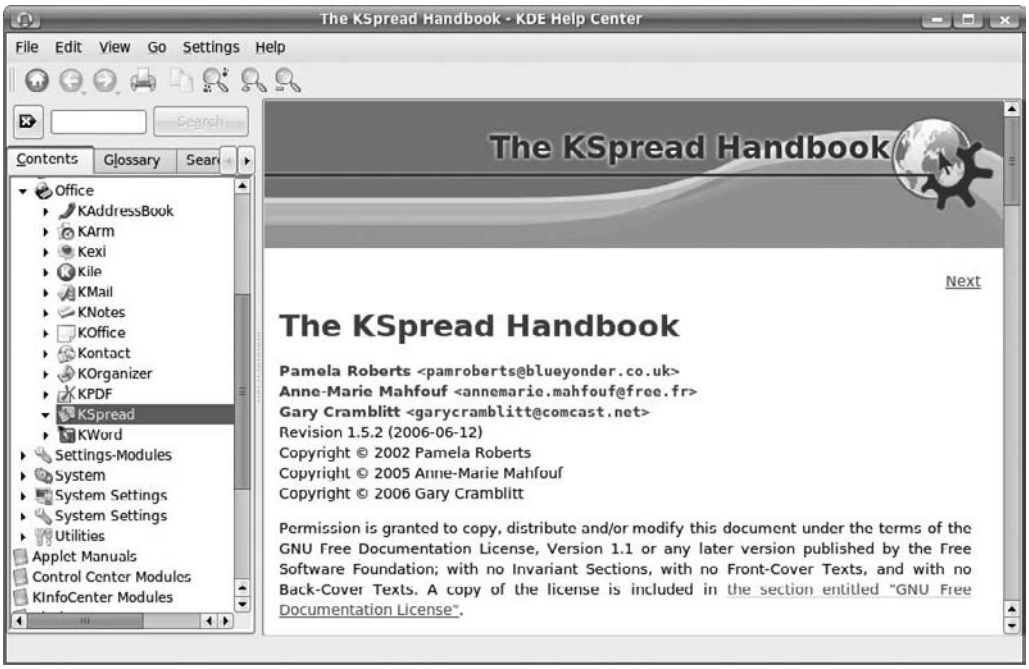


Any cell in KSpread that contains a calculated value displays a blue triangle in the lower-left corner of that cell. You can see this in Figure 17-35, where all of the cells in column E contain calculated values.

KSpread provides extensive online help, available by selecting Contents from the Help menu, or by pressing the F1 key on your keyboard at any time. Figure 17-36 shows the main page for KSpread's online help, known as the KSpread Handbook. This is actually a fairly long and quite complete document. To see the table of contents, scroll down in the right-hand portion of the window.

FIGURE 17-36

KSpread's online help, The KSpread Handbook



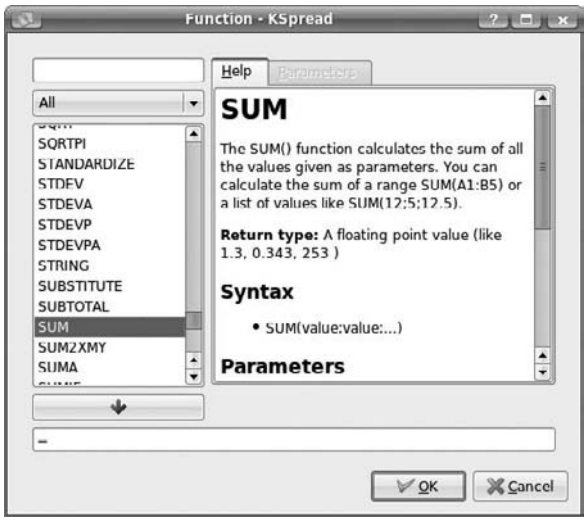
Invoking Functions in KSpread

The tutorial given earlier in this chapter explained how to manually enter a function, which is fine if you are using a simple function that you happen to know the name of. KSpread supports an extensive collection of functions that you can easily select from the Insert ➪ Function dialog using the mouse. Figure 17-37 shows this dialog with the sum function selected from the complete list of available formulas.

As you can see, selecting any function from the Function dialog displays fairly extensive help for that function in the pane at the right of the dialog — in this case, "The SUM() function calculates the sum . . ."

FIGURE 17-37

Available functions in KSpread



Specifying the Type of Data in a Cell

The spreadsheet tutorial at the beginning of this chapter used fairly straightforward types of data that KSpread could easily recognize — sequences of characters from the alphabet and spaces are probably words, whereas numbers and decimal points are probably numeric values. Knowing the type of data that is contained in a cell (generally referred to as its data type or data format) is fairly important, especially when you subsequently want to invoke functions that calculate values from the contents of various cells.

When you enter values in a spreadsheet without specifying their type, spreadsheets such as KSpread assign them a general data format, which performs a best guess of their data format based on the type of characters that they contain. However, you can also manually specify the format of any cell (or cells) by selecting those cells and using the commands on the Format ➪ Cell Format dialog, shown in Figure 17-38.

TIP **You can also display the Cell Formats dialog at any time by right-clicking on a cell and selecting the Cell Formats menu command from the context-sensitive menu, or by selecting a cell and pressing Alt+Ctrl+F.**

As an example, note that cell C3 in the example spreadsheet (most recently shown in Figure 17-35) doesn't display the traditional number of decimal places that you'd expect to see in a monetary value. This is because KSpread's general data format recognized it as a decimal number, but didn't have any way of knowing that this was a currency value. To correct this, select that cell, select the Format ➪ Cell Format command, and click Money from the list at the left. This displays a list of possible representations for that data format, as shown in Figure 17-39.

FIGURE 17-38

Specifying the type of data in a cell in KSpread

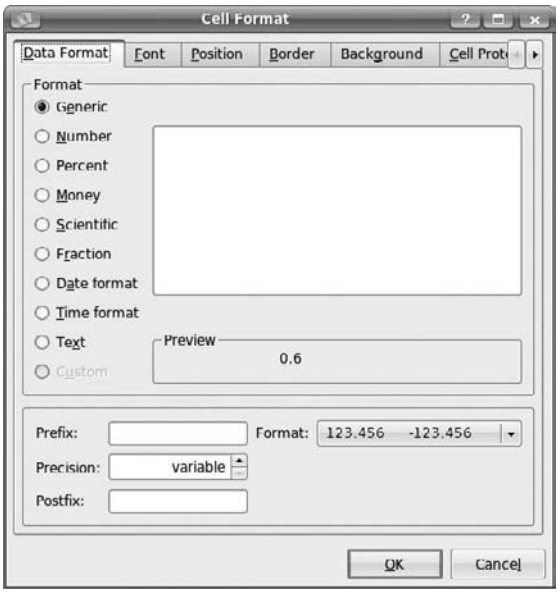
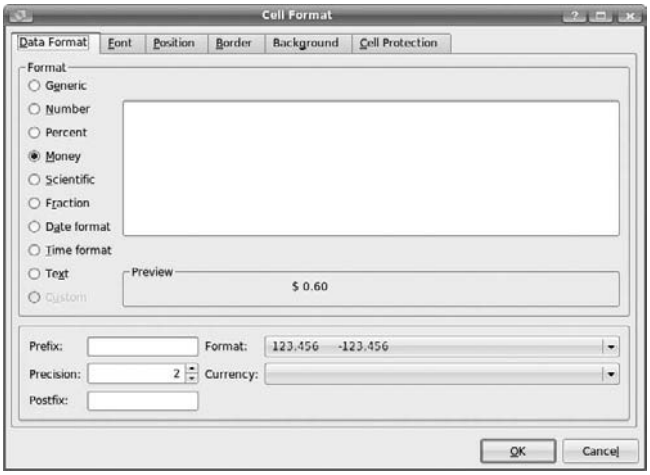


FIGURE 17-39

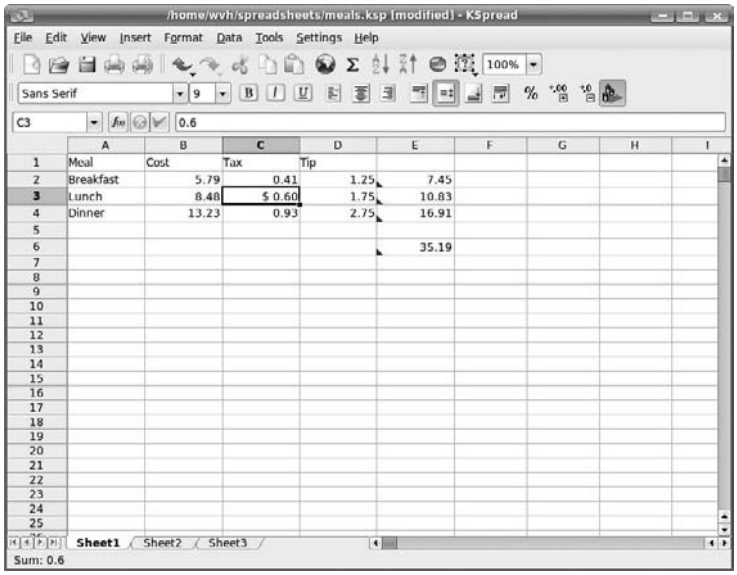
Specifying that a cell contains a currency value



Note that the Precision field (at the lower left) specifies that two decimal places is the default for this data type, which is what you are trying to correct. Because this would do the trick, click OK to set cell C3 as a currency value and close the Cell Format dialog. This updates the spreadsheet, which now shows the correct number of decimal places for cell C3, as shown in Figure 17-40.

FIGURE 17-40

The updated spreadsheet showing C3 as a currency value



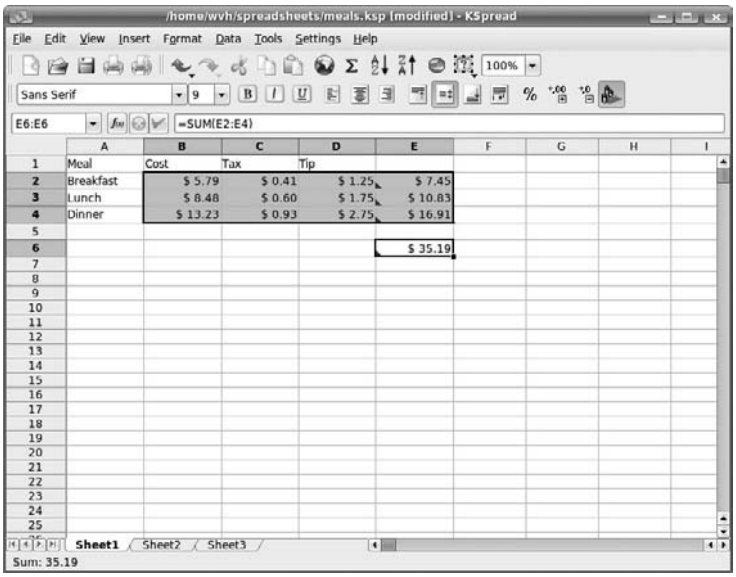
This is fine, except that KSpread now recognizes that cell as a currency value but all other numeric cells are still general numeric values, which makes the spreadsheet look a bit odd. You can correct this by multi-selecting all of the numeric values in the spreadsheet (hold down the Ctrl key, click and hold the left mouse button on cell B2, drag the selection box to encompass all cells from B2 through E4, and then also click cell E6 before releasing the Ctrl key), displaying the Format ➪ Cell Formats dialog again, selecting Money, and clicking OK. Voilà! The spreadsheet is now consistent, as shown in Figure 17-41.

Importing Existing Spreadsheets into KSpread

Linux folks like me tend to be Linux-centric, but that doesn't mean that we can completely ignore reality. When you're working with existing spreadsheets (and any other office-related documents, to be honest), the chances are pretty good that you created them using more common software packages, such as Microsoft Excel, that run on other operating systems, such as Mac OS X and Microsoft Windows. The folks who wrote and support KSpread couldn't ignore this reality either, so KSpread provides built-in support for opening Microsoft Excel spreadsheets and workbooks.

FIGURE 17-41

The updated spreadsheet showing currency values



NOTE **If you have a lot of existing spreadsheets and are wondering how to transfer them to your new Ubuntu system, see Bonus Chapter 2, "File Transfer and Sharing on Ubuntu," for detailed information on doing just that. Don't worry — it's really quite easy to do using a network connection or even a traditional mechanism such as a floppy (if you still use them), a flash drive, or a CD.**

For example, to open an existing Excel spreadsheet, simply select the File ➪ Open command. By default, the file selection dialog that opens enables you to select files with any extension, including the `.xls` extension used by Microsoft Excel spreadsheets.

The same mechanism is used to open files in any other format supported by KSpread. KSpread can open spreadsheets or other data files that have been saved in any of the following formats:

- **Applix** (`.as`)**:** The spreadsheet files produced by the Applix Office suite, which was discussed in Chapter 16.
- **dBASE Document** (`.dbf`)**:** The database format used by programs such as dBase III, old versions of FoxPro, and most applications that use the XBase format. XBase is the generic way of referring to applications that use this file format.
- **GNUmeric** (`.gnumeric`)**:** The XML file format used the Gnumeric spreadsheet applications, which was discussed earlier in this chapter.
- **Kexi Database File-based Project** (`*.kexi`)**:** The SQLite database file format used by Kexi, an open-source application that is designed for compatibility with Microsoft Access. For more information about Kexi, see `www.kexi-project.org`. For more information about the SQLite database, see `www.sqlite.org`.

- **Quattro Pro Document** (`.wb1`, `.wb2`, `.wb3`)**:** The spreadsheet formats used by various versions of Borland's Quattro Pro spreadsheet.
- **OASIS OpenDocument Spreadsheet** (`.ods`)**:** The open document spreadsheet format used by OpenOffice.org Calc and other forward-thinking (i.e., open) spreadsheets.
- **Plain Text Document** (`*.txt`)**:** A text file containing rows and columns of data, delimited by whitespace.
- **Text File with Comma-Separated Values** (`.csv`)**:** A standard text-based interchange format, where each row appears on a single line and each field value is separated from the next by a comma. This exchange format does not preserve the calculations required to produce cell values, only the values of the cells themselves.

As you can see, the open source developers who are responsible for KSpread have made a very significant effort to ensure that you can import any of your existing spreadsheet data files into KSpread, so that you're not losing anything other than your dependence on your previous operating system.

Creating and Using OpenOffice.org Presentations

The OpenOffice.org Office suite includes a powerful software package known as Impress for creating and editing presentations. As a tip of the hat to usability at the expense of confusion, the menu item for invoking this piece of software is labeled OpenOffice.org Presentation, which is what it creates. While you may not need to create presentations on your home computer, presentations are a fact of everyday life in business today, and no suite of office applications would be complete without this type of package. As mentioned earlier, I'll simply refer to OpenOffice.org Impress as Impress throughout the remainder of this section.

Installing Files for Impress

The OpenOffice.org software, including Impress, is installed by default if you install Ubuntu from the Live CD, from the Alternate Install CD using the Text Mode installer, or from the Alternate installer using the OEM Mode installer, or if you install Kubuntu from its Live CD. If you want to add Impress or the entire OpenOffice.org suite to a server system that has a graphical user interface, you can install it using apt-get, aptitude, Synaptic, or the Adept tools, as explained in Chapter 23, "Adding, Removing, and Updating Software."

One of the things that differentiates Ubuntu and Kubuntu Linux from other Linux distributions is their extreme dedication to the user, especially in terms of internationalization. If you are installing Impress for use when creating presentations in languages other than English, you may also want to search for and install the appropriate hyphenation, localization, and thesaurus packages for whatever other locales you are creating spreadsheets for.

NOTE **If you've already installed the localization and internationalization packages for any component of the OpenOffice.org suite, there is no need to reinstall them. Installing them once installs them as a shared resource that can be used by every component of the suite.**

You can easily find these localization and internationalization packages by looking for a string that represents the country code that you are looking for. For example, Figure 17-13 (earlier in this chapter) showed the specific German language support packages for OpenOffice.org that have been identified by searching for the string "-de."

TIP **You can install complete language support for all applications (without translating the GNOME user interface itself) by installing the appropriate language support meta-package for your target locale(s). These packages have names of the form language-support-country-code (for example, language-support-de), and include the dictionaries for the specified locale. To fully translate the GNOME interface and all internationalized applications, select the appropriate language pack. The names of the language packs are of the form language-pack-country-code. Once you've installed multiple dictionaries, you can set the default language for your document using the Tools ➪ Options ➪ Language Settings ➪ Language drop-down list.**

See Chapter 23, "Adding, Removing, and Updating Software," for more detailed information on installing applications using the Synaptic Package Manager or the Adept tools.

Starting Impress

You can start Impress from any command-line by executing the command `ooimpress`, but it is usually easier to start it from your GNOME menus. A menu item for Impress is provided in the Applications ➪ Office menu on Ubuntu systems and on the K Menu ➪ Office menu on Kubuntu systems. To start Impress from these menus, select the OpenOffice.org Presentation menu item.

Using the Presentation Wizard

When Impress starts, it runs a Wizard that walks you through creating a presentation. Let's follow along.

Figure 17-42 shows the initial screen of the Impress Presentation Wizard.

FIGURE 17-42

The first screen of the Impress Presentation Wizard



Assuming that you haven't opted out of the Presentation Wizard, click Next to begin creating your new presentation. The second Presentation Wizard screen displays, as shown in Figure 17-43.

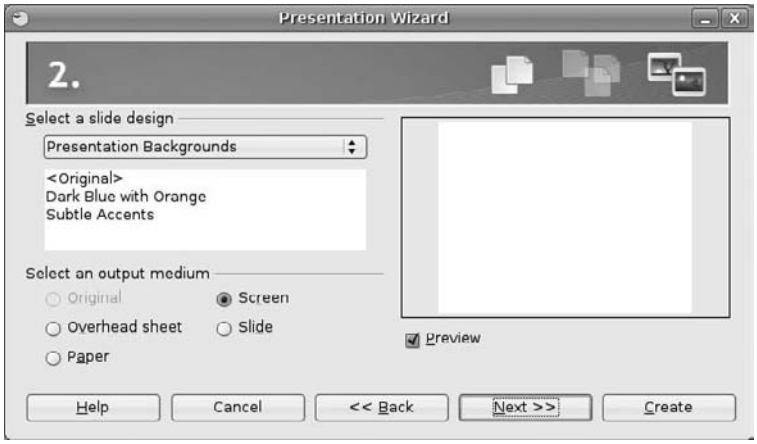
TIP **When you start Impress and see the Presentation Wizard, you can opt out of the Presentation Wizard by selecting the Open existing presentation radio button. Selecting this button enables you to select an existing Impress or Microsoft PowerPoint presentation using the Open dialog that displays. If the presentation that you want to open is located in your home directory, you can double-click its name to open it. If it is located in another directory, click Open to display a dialog that enables you to browse for the presentation that you want to open.**

You can permanently opt out of starting Impress with the wizard in the future by clicking Do not show this wizard again before you proceed. Impress will remember this setting and display a blank presentation the next time that you start Impress.

If you do not have any existing presentations that you want to open and simply want to open a new, blank presentation without using the wizard, click the Create button while the Empty presentation radio button is selected.

FIGURE 17-43

Specifying a background in the Presentation Wizard



Select one of the available backgrounds for your presentation from the list (or keep the <Original> setting to continue without a background) and identify the output medium that you are targeting for this presentation: an onscreen presentation, overhead sheet, paper, photographic slides, and so on. Specifying the output medium modifies some internal settings to maximize the appearance and usable area in your presentation. Once you've made any desired changes to these settings (I selected the Dark Blue with Orange presentation background for this example), click Next to continue. The third Presentation Wizard dialog displays, as shown in Figure 17-44.

This screen enables you to select an effect that Impress should use when changing from one slide to the next. These include many different sweeping movements, rotations, and so on. Rather than try to describe them all here, I'll just suggest that you experiment with them. Selecting any effect from the list while the Preview checkbox is selected (which it is by default) shows a preview of the selected effect in the preview window on this screen of the wizard.

FIGURE 17-44

Providing presentation details in the Presentation Wizard

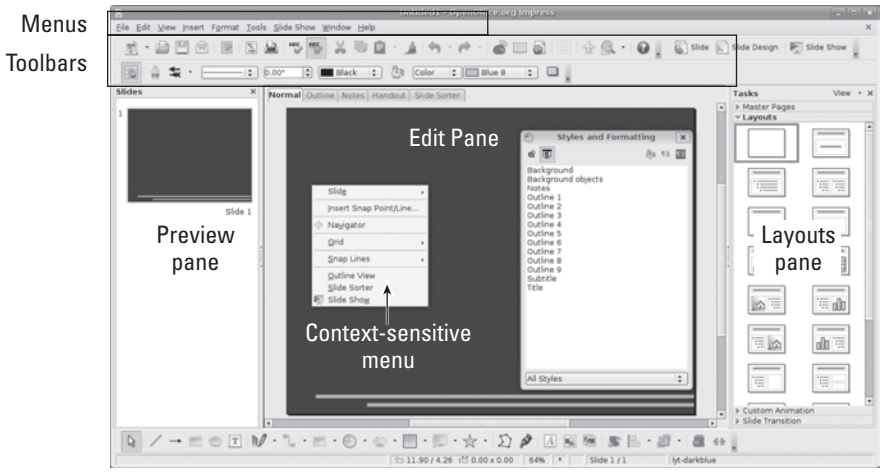


The other values on this wizard screen enable you to specify the speed of any selected effect, and the type of presentation that you are creating. Default presentations change slides only when you press a key or otherwise indicate that they should change. Automatic presentations change slides after a specified period of time, and are intended for automated use in kiosks, perhaps at trade shows where one might want a regular, possibly repeating, series of slides to serve as an attractor.

Once you've changed any of the settings on this screen that you want to take advantage of, click Create to create your new presentation and complete the Presentation Wizard. The standard Impress screen displays, as shown in Figure 17-45.

FIGURE 17-45

The standard Impress screen displaying a new presentation



Taking a Quick Tour of Impress

Impress provides excellent online help, as do all of the tools in the OpenOffice.org suite, so I won't bore you by walking you through each menu and entry that it contains. However, as a quick overview of the Impress interface, the callouts in Figure 17-45 highlight the following areas:

- **Menus:** The drop-down menus in Impress are organized much like the menus in Microsoft PowerPoint to provide a familiar user experience. Menu items with an arrow at their right edge lead to other, related submenus.
- **Toolbars:** By default, Impress displays the Standard, Presentation, Drawing, and Line and Filling toolbars, which provide one-click access to specific commands when you click on the appropriate icon. You can modify the list of icons displayed on a toolbar by clicking on the control for that toolbar, located at the far right. You can display other toolbars or deactivate the default toolbars by selecting the names of those toolbars from the View ➪ Toolbars submenu.
- **Preview pane:** By default, the pane at the left displays a preview of all of the slides in the current presentation.
- **Edit pane:** By default, the central pane displays your working copy of the current slide. You can change the way in which this slide is displayed by selecting the other tabs at the top of this window.
- **Layouts pane:** By default, the pane at the right displays the predefined layouts that can be applied to the current slide. You can change the contents of this window by selecting other presentation-related tabs, such as Templates, Master Pages, Custom Animations, and Slide Transitions.
- **Context-sensitive menu:** Right-clicking on the slide that you are working on in Impress displays the context-sensitive menu shown in Figure 17-45. This menu provides quick access to several cell manipulation commands.

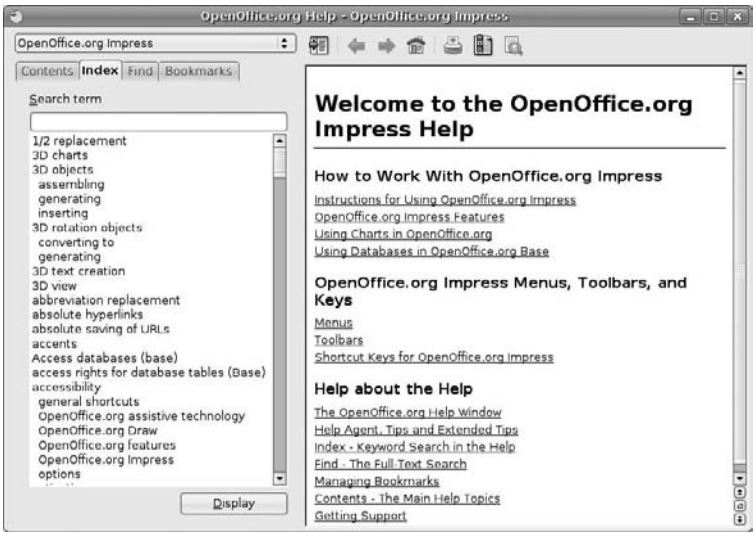
TIP **Although it's not displayed by default, you can display the Styles and Formatting window at any time by pressing the F11 key on your keyboard. (Press F11 again to close this window.) The Styles and Formatting window provides easy access to the styles that are available in the current presentation, and is especially convenient if you choose not to display the formatting toolbar. Figure 17-45 shows Impress with the Styles and Formatting window displayed. Like most floating windows in OpenOffice.org applications, you can dock this window by holding down the Ctrl key and double-clicking on the gray portion of the window heading, making it a part of the default OpenOffice.org interface rather than a floating window. You can undock the window by clicking and dragging on the gray portion of the window heading, or by holding down the Ctrl key and double-clicking on the heading of the docked area.**

Impress provides extensive online help, available by selecting OpenOffice.org Help from the Help menu, or by pressing the F1 key on your keyboard at any time. Figure 17-46 shows the main page for Impress help, which provides both a searchable index and a list of popular, general topics from its initial display window (at right).

TIP **If you are having problems viewing the text for the online help in Impress, and OpenOffice.org Writer is not installed on your system for some reason, try installing OpenOffice.org Writer. This will force an update to the components that control online Help in the OpenOffice.org suite and will usually resolve the problem.**

FIGURE 17-46

Online help for Impress



Creating a Presentation

Creating presentations in Impress is easy, although fine-tuning them can be as challenging as it is in any software for creating presentations.

To apply any of the available formats to the current slide, select the thumbnail for that format from the Layouts window. The selected layout is applied to the current slide. You can then click in any of the text boxes and enter the appropriate text for that region. Figure 17-47 shows a slide that uses the Title Slide layout at the upper right in the Layouts window.

Impress provides many sample layouts, including title slides, slides that provide headings and text areas, slides that provide headings and areas for graphics, and so on. You can identify the intent behind any of the available layouts by positioning the cursor over the thumbnail for that layout and viewing the tooltip that displays.

Once you've entered your text or other content in the current slide, you can add a new slide by selecting the Insert ➪ Slide command. A new slide displays in the editing window, and the preview window is updated to display the new slide and any previous slides. You can then select the thumbnail for the format that you want to use on the new slide. For example, Figure 17-48 shows a new slide with the Title, two Text Blocks format applied, and a preview of the previous and current slides.

Once you've added a number of slides to your presentation, you should save it using the File ➪ Save command. If this is a new presentation, you will be prompted to supply a name for the presentation.

After saving your presentation, you will probably want to see what it looks like so far. To preview your actual presentation, select View ➪ Slide Show or press the F5 key on your keyboard. Your presentation displays on your screen, using any transition effects and timing that you specified when you created the presentation.

FIGURE 17-47

Creating a title slide in Impress

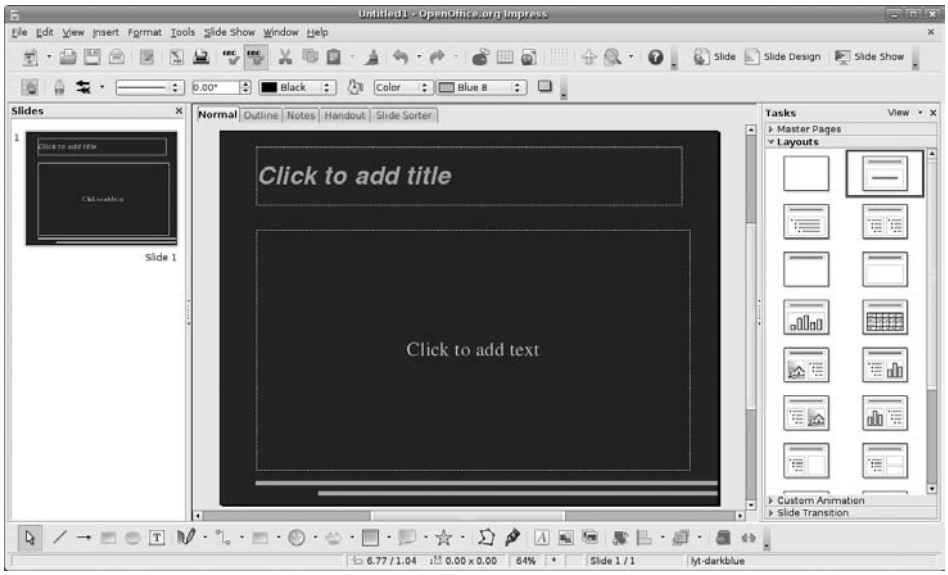


FIGURE 17-48

Adding a new slide in Impress



TIP **When you preview your presentation, the preview starts at the first slide selected in the preview pane. To see your entire presentation, select the first slide in the presentation before requesting a preview.**

Importing Existing Presentations

Impress can open existing presentations produced by Microsoft PowerPoint (and those produced by itself or StarOffice, of course). After working on these, you can then save them directly in the new, open OpenDocument presentation format that is used by Impress.

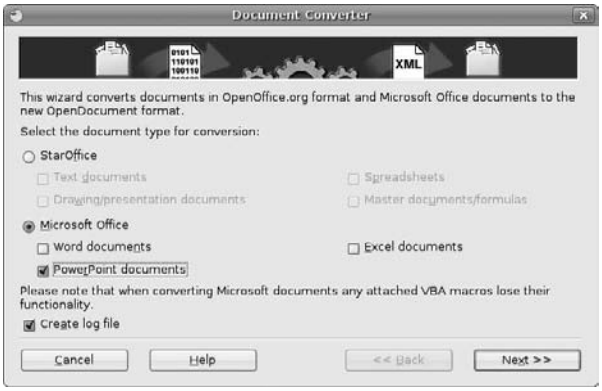
Repeating this process is fine if you have only one or two presentations that you want to update, but if you are moving all of your work to Ubuntu, you may find that you have many presentations that you want to update to Impress format. Fortunately, Impress provides a wizard for this that will not only import all of your Microsoft PowerPoint presentations, but also will import and convert any existing or attached layout files and master documents that it detects. This wizard is known as the Document Converter. You can start this wizard by selecting the File ➪ Wizards ➪ Document Converter menu item.

TIP **At the time of this writing, the Java GUI widgets used by some of the OpenOffice.org wizards (such as the Document Converter) have problems with the compositing effects used by Compiz on Ubuntu desktop systems. If you are trying to run an OpenOffice.org wizard and it looks like a really whizzy postage stamp that can't be resized, exit the OpenOffice.org app and use Ubuntu's Preferences ➪ Appearance app to temporarily disable advanced visual effects by selecting None from the Visual Effects tab. You'll then be able to use the wizard. You can always re-enable the special effects once your alien presentations are all in OpenOffice.org's ODP format.**

The Document Converter used to convert presentations works identically to the Document Converter used to batch import spreadsheets, as explained earlier in this chapter, so I won't waste space by repeating those instructions here. For information on using the Document Converter Wizard to batch-import documents of any supported type, see "Importing Spreadsheets Using the Document Converter" earlier in this chapter. The only real difference when importing presentations rather than spreadsheets (or word processing documents, as explained in Chapter 16) is that you must select Microsoft Office and PowerPoint documents on the initial screen of the wizard, as shown in Figure 17-49.

FIGURE 17-49

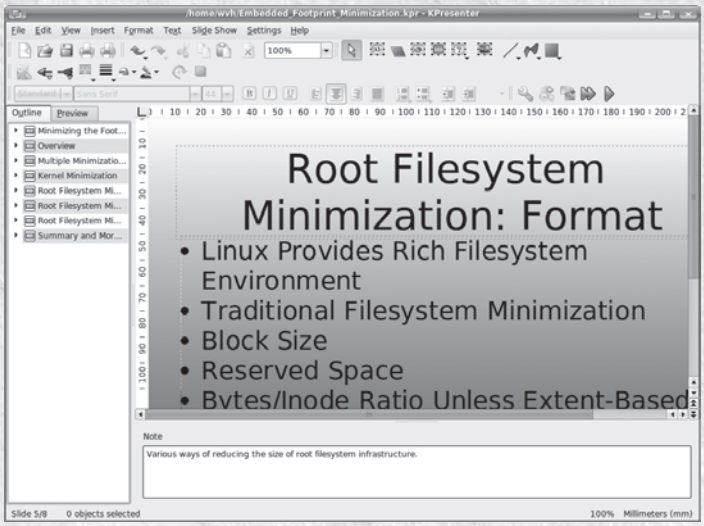
Specifying the input format when converting presentations



As shown in Figure 17-49, converting any Microsoft documents will make "any attached VBA macros lose their functionality." And that's a bad thing? You might also say that converting any Microsoft documents will cause them to lose much of their ability to carry and launch viruses. This seems like a selling point for using OpenOffice.org on all platforms, but your mileage may vary.

Hey, What About KPresenter?

KOffice fans may feel that I've been remiss by not devoting significant space to KPresenter, the presentation software component of the KOffice suite. KPresenter is not installed by default on Kubuntu systems, but can be added by using any of the software installation tools discussed in Chapter 23 to install the kpresenter and kpresenter-data packages on Kubuntu systems running the KDE 3.5.*x* system, or the kpresenter-kde4 package to your Kubuntu system running KDE 4.*x*. You can also install them on your Ubuntu system, for that matter. As you can see from the accompanying figure, KPresenter is a nice package that does a fine job of creating presentations — I just haven't used it extensively and don't feel that I know it well enough to discuss it in depth.



A presentation in KOffice KPresenter

If you're a KOffice fan, KPresenter is worth checking out for familiar reasons: It is lighter weight than the OpenOffice.org tools, provides the same interoperability with the other KOffice components that those other components do, and was designed for use on KDE-based systems. KPresenter can import OpenOffice.org, Microsoft PowerPoint, and XML files, and presentations done in the format used by the old FSF MagicPoint application (`http://member.wide.ad.jp/wg/mgp/`).

Summary

This chapter discussed common types of office software that you may need to use to complete your move to Ubuntu or Kubuntu Linux. Linux often gets a bad rap in the press for not having a rich set of application software — as this books shows, this is entirely false. Not only is a good selection of powerful office software available for Linux, but the open source nature of Linux gives you many alternatives in terms of the package that you want to use — and all at the same low, low price, which is $0 thanks to the hard work and time spent by the open source community and the folks who bring you Ubuntu and Kubuntu Linux.

This chapter began with an introduction to spreadsheets for readers who haven't had the need or the opportunity to use them before. The next section discussed OpenOffice.org Calc, the powerful spreadsheet component of the OpenOffice.org Office suite that is installed by default on both Ubuntu and Kubuntu desktop systems. The following two sections discussed other popular spreadsheets for Ubuntu and Kubuntu Linux, Gnumeric and KSpread, respectively. As you were able to see, all of these are basically similar, but each provides some "special sauce" that might make you prefer it to the others. The chapter concluded by discussing OpenOffice.org Impress, which is a powerful software package for creating, working with, and managing presentations, and is another part of the OpenOffice.org suite of office applications.

Now that we've explored creating documents, spreadsheets, and presentations, the next chapter discusses the software that enables you to create and convert various types of graphics files on Ubuntu and Kubuntu desktop systems.

Chapter 18

Working with Graphics

As a writer, I've always somewhat resented the old saying "A picture is worth a thousand words." Can't I just do the thousand words and we can all move on?

I'm kidding, of course — there are many times when a picture is worth far more than any number of words, such as the screenshots in this book, photographs and artwork in advertising, or various publications that specialize in celebrating the human form. No one has ever described a text-heavy document as eye-catching, despite even the best efforts of typographers and the layout artists at *Wired*. Even if you're not a graphics artist, there are times when everyone needs software for editing graphics and creating drawings, whether you just want to remove the terminal case of red-eye from all of your vacation photographs or produce simple illustrations for a card, letter, or presentation.

If you've worked with graphics packages such as Adobe Photoshop, Adobe Illustrator, CorelDRAW, or Corel's Paint Shop Pro (formerly from JASC), you'll be happy to know that open source packages with the capabilities that you'll need to do the same sort of work are freely available for Ubuntu and Kubuntu Linux. Some are pre-installed, and some are located in Ubuntu repositories and are therefore just a download away.

This chapter helps you make some sense out of the maze of abbreviations, acronyms, general terminology, and specialized jargon that aficionados use to describe digital graphics, and explains how to use the software available for Ubuntu and Kubuntu Linux to create and edit drawings, images, and other types of computer graphics.

IN THIS CHAPTER

- Overview of digital graphics
- Working with images in GIMP
- Photo editing overview
- Using Krita
- Using OpenOffice.org Draw
- Working with vector graphics

TIP **If you are creating graphics and need a starting point or some inspiration, you'll be happy to know that an open source library of clip art is available to complement the open source software that is discussed in this chapter. The aptly named Open Clip Art Library (`http://openclipart.org`) is a collection of user-contributed clip art that has been placed in the public domain according to the Creative Commons license for that purpose (`http://creativecommons.org/licenses/publicdomain`). To quote the license, this allows the graphics in this library to be "freely reproduced, distributed, transmitted, used, modified, built upon, or otherwise exploited by anyone for any purpose, commercial or non-commercial, and in any way, including by methods that have not yet been invented or conceived." There's a lot of nice stuff there — check it out!**

Overview of Digital Graphics Terminology

Even if you're not a graphics artist (and I certainly am not one), everyone with a Web site or a digital camera has some familiarity with digital graphics. Terminology-wise, digital graphics is basically nothing more than storing images in files on your computer, regardless of whether these files contain pictures from your summer vacation, fancy buttons or other images that you're using on your Web site, artwork and drawings that you've created using graphics software, pictures that you've used a scanner to capture, images captured directly from your computer screen, or just about anything else that you can store in a file and display in a graphics application.

The following provides a quick glossary of graphics terms that you may encounter as you start to work with online, digital graphics:

- **Anti-aliasing:** A technique used to reduce the jagged appearance of the edges of objects and text in graphical images. Anti-aliasing softens those edges by blending them with the colors or areas that appear beside them in an image.
- **Bitmap:** A graphical image that is composed of multiple picture elements, each representing one bit of that image. Each bit in a bitmap has its own color value, generally expressed as a combination of red, green, and blue (RGB) color values. Bitmaps are one of the two main formats for digital images — the other is vector graphics. Bitmaps are also referred to as raster images. The generic bitmap should not be confused with the common graphics format known as bitmap, or BMP.
- **BMP:** A graphics format originally developed by Microsoft that supports a variety of resolutions and color depths. A file in BMP format contains a single bitmap/raster image.
- **CMYK:** A color definition model that identifies the amount of **C**yan, **M**agenta, **Y**ellow, and Blac**K** in a given color. CMYK color designations are commonly used in printing because these are the four traditional ink colors used in commercial printing. Cyan, magenta, and yellow are referred to as *subtractive primary colors* because they are based on reflected light and thus begin with white light and subtract wavelengths to display specific colors. In the commercial (paper) print process, a separate printing plate is made for each of the CMYK colors. These are therefore often referred to as *process colors*.
- **Color depth:** The maximum number of colors contained in an image or that can be displayed by a given device or graphics card. Common color depths are 256 colors, 16-bit (65,535 colors), 32-bit (4,294,967,296 colors, generally simply referred to as "millions of colors"), and so on. The greater the color depth of an image, the greater its size, because the color of each pixel requires the associated number of bits to represent each of the red, green, and blue portions of its color. Color depth is often defined in "bits per pixel" (bpp). For example, high color uses 15 or 16 bpp: 5 bits for blue, 5 bits for red, and 5 or 6 bits for green.

- **Color space:** A model for representing color. RGB, CMYK, and HSV are common color spaces.
- **Dithering:** The process of displaying pixels of different colors in a portion of an image to give the visual impression of a single color. Dithering is often used to give a lower-resolution image the smoother appearance of a higher-resolution image.
- **GIF:** A popular lossless graphics format, Graphics Interchange Format files were formerly quite popular on the Web because they produced relatively high-quality, small file size images through the use of a compression technique known as Lempel-Ziv-Welch (LZW) compression, which is patented by Unisys. Because of licensing issues, most Web sites now use JPEG or PNG graphics. The primary advantage of the format is that a single GIF file can contain multiple images that can be configured to display in sequence. These multi-image GIF files are known as "animated GIFs." To avoid patent and licensing issues, most animated Web graphics today are created using file formats such as Flash, QuickTime, or Shockwave, or are created using Java-based animation techniques.
- **Grayscale:** An image that is composed of various shades of gray. The number of bits used to hold the color value for each pixel determines the number of shades of gray that the image can use. If the number of bits is 1, then the image is simply black and white, which is also known as *monochrome*. Grayscale is typically used in newspapers or books that are printed in black and white, where a halftone is used to print the different shades of gray.
- **Halftone:** The simulation of a continuous-tone image (shaded drawing, photograph) by using different-sized dots to make up different portions of the image. Smaller dots per print area are used for lighter areas, and larger dots are used for darker areas. Halftones are typically used for relatively high-resolution reproduction of black and white images in print, such as in a newspaper or magazine.
- **HSV:** A color definition model in which colors are represented by their hue (X axis), saturation (Y axis), and brightness (Z axis).
- **Jaggies:** The term for the staircase effect seen on the edges of bitmapped objects, also sometimes referred to as "aliasing." This unpleasant visual effect can be mitigated by the use of techniques such as anti-aliasing.
- **JPEG:** A popular bitmapped image format consisting of 24-bit color and developed by the Joint Photographic Experts Group. JPEG is a lossy format, which is often used for graphics on Web pages to provide a good balance of good image resolution and small file size.
- **Layers:** Most complex graphics files are internally composed of multiple image layers that are superimposed to create the final image. Working with different portions of an image (background, each component of an image, and so on) not only makes it easy to interact with each graphics element without disturbing others but also makes it easy to quickly build up complex graphics by separating them into different logical elements that can be reused across similar graphics.
- **Lossless formats:** File formats that are not lossy, i.e., file formats that provide all of the information that was present in the audio, graphics, or video information as it was originally recorded, drawn, or captured.
- **Lossy formats:** File formats that sacrifice image/audio quality to reduce file size. The term *lossy* refers to the data compression algorithm used to reduce file size, where some of the information that was present in the original file is actually eliminated to produce the smaller, compressed file, and therefore the original file can never be completely reconstructed from the compressed version. In the case of graphics, lossy images typically reduce both the available palette of colors and the number of pixels in the image. Lossy formats are quite common in graphics, audio, and video files because they can often reduce file size by eliminating information that is beyond the range of the human eye, the human ear (e.g., mp3 files), or both.

- **Metadata:** Data about other data. In graphics files, metadata typically provides information about file data elements, file attributes (name, size, location, color depth, data type, and so on), or free-form comments. Metadata simplifies the use and analysis of graphics files by providing a higher-level interface to the file than just the raw data itself.
- **Monochrome:** Images and text that are displayed using only black and white pixels. The term comes from the fact that there is only one chromatic (i.e., color) value used in the image: black is a pixel that is set to that color value, while white is a pixel that is not set to black.
- **Pantone colors:** A set of specific color designations created by Pantone, Inc. (`www.pantone.com`) that are used to provide definitive color values for various portions of an image. The set of these standardized colors is known as the Pantone Matching System (PMS). Pantone colors are identified by unique identification numbers in a color guide known as a swatch book. Pantone colors are used to identify spot colors in an image — in other words, portions of an image that do not blend with any other portions of the image.
- **PCL:** A Printer Control Language originally developed by Hewlett-Packard (HP) and now used in thousands of different printers. Files in PCL format can be printed directly on HP and compatible printers with no preprocessing.
- **Pixels:** The smallest graphical element in a digital image that can be displayed on your computer screen. Each pixel has its own color, generally expressed as a combination of red, green, and blue (RGB) color values. The term *pixel* stands for "picture element." Most graphical images, such as digital photographs, contain millions of pixels — hence the term megapixels. The more pixels an image contains, the more detail it provides and the higher its resolution.
- **PNG:** Portable Network Graphics, a graphics file format originally designed as an alternative to GIF and its attendant licensing problems. PNG files are lossless, supporting up to 48-bit color or 16-bit grayscale images without losing any detail during compression.
- **PostScript:** A printer control language developed by Adobe and used in thousands of high-quality laser printers. PostScript files support both bitmaps and drawing primitives (low-level commands that create graphics by drawing lines from point to point, generating geometric shapes using corner, center, perimeter coordinates and so on), and therefore typically produce high-quality images. PostScript was inspired by an earlier printer control language called InterPress, which was developed at Xerox PARC, as were other handy things such as Ethernet, laser printers, bitmapped displays, and so on.
- **Prepress:** The process of preparing documents for printing. Nowadays, this generally refers to the use of computers to prepare documents and illustrations for printing as opposed to the use of analog techniques in the past for such activities as the production of photographic negatives, manual page layout, manual typesetting, and the manual production of printing plates.
- **Raster graphics:** Another name for bitmap graphics. The term *raster* refers to a position in the grid of coordinates that make up a computer display.
- **Resolution:** The output quality of an image, defined in terms of dots, lines, pixels, or samples per inch, depending on your output device. PPI (pixels per inch) refers to screen resolution, DPI (dots per inch) refers to printer resolution, SPI (samples per inch) refers to the resolution at which images are obtained from a scanner, and LPI (lines per inch) refers to halftone resolution. High-resolution images contain 300 or more units per inch, and are intended for printed output. Low-resolution images generally contain 100 or fewer units per inch and are generally intended for onscreen or Web-based display.

- **RGB**: A color definition model that identifies the amount of **R**ed, **G**reen, and **B**lue in a given color, typically represented as either three numbers from 0 to 255 or a series of six hexadecimal digits in which each pair of digits represents the intensity of each color. For example, 255 0 255 and FF00FF represent the same color, a bright purple comprising the maximum value of red, no green, and the maximum value of blue. All computer monitors use RGB colors. Red, green, and blue are added together to produce white, and are therefore referred to as *additive primary colors*.
- **Stair-stepping:** Another term for jaggies.
- **TIFF:** Tagged Image File Format, a popular lossless graphics file format. Files in TIFF format often have the extensions .TIF or .TIFF. TIFF files are typically quite large and therefore are unsuited for display on Web pages.
- **Vector graphics:** A digital graphics representation and associated file format that defines an image by identifying each object in that image and describing that object in terms of characteristics such as its outline, shape, color palette, associated text, gradients, patterns, and element grouping. Vector graphics are resolution independent because they can be redrawn at the highest quality supported by a graphics application or display device.

For a general source of information about digital graphics, the best site I've found on the Web is About.com's Graphics Software section at `http://graphicssoft.about.com`. This site contains everything from a great dictionary of graphics terms to hundreds of useful tutorials, articles, and discussions. It's truly amazing! The only problem with this site is that it is titled, "About Graphics Software for PC or Mac." Hey, what about the rest of us? Luckily, as you'll see throughout this chapter, the graphics tools available for free on your Ubuntu and Kubuntu systems are more than sufficient for almost every graphics task.

Using GIMP

GIMP stands for GNU Image Manipulation program, the premiere open source graphics package for Linux systems. GIMP originally stood for General Image Manipulation Program and was created by Spencer Kimball and Peter Mattis. From the beginning, GIMP was designed to be easily extendible by the addition of plug-ins that could be written, compiled and maintained separately, but which would work seamlessly within the GIMP framework. Like any other open source package, GIMP has substantially benefited from the contributions of others, especially in terms of plug-in development. GIMP is a shining example of a powerful and successful open source package, having also been ported to Microsoft Windows and Apple's Mac OS X and used extensively on those platforms. For all of the details of GIMP history, see `www.gimp.org/about/ancient_history.html`.

GIMP is also legendary among GNU software packages, especially as far as the GNOME desktop environment goes, because the toolkit on which much of the GNOME user interface is based is GTK, the GIMP Tool Kit. Without restating everything on the GIMP ancient history page, let's just say that GIMP started out using another X Window system toolkit, but its authors decided to create their own, which was full-featured enough to be adopted by many other projects. Biblically, GIMP begat GTK, which begat GTK+, which begat GNOME, which now supports GIMP. I hope that's clear — there will be a quiz later.

NOTE **People often refer to GIMP simply as gimp or the gimp (not to be confused with the character from Quentin Tarantino's *Pulp Fiction*), because that's the name of the command that you type at the command-line to start the program. Throughout this chapter, I'll use GIMP as a way of generally referring to the program, using `gimp` (lowercase) only when I'm specifically referring to the command-line command.**

Starting GIMP

GIMP is automatically installed as part of a basic Ubuntu desktop system install, so no special installation process is required on Ubuntu systems. On Kubuntu systems, because of its GNOME roots and reliance on GTK, you'll have to use the familiar apt-get, aptitude, or Adept tools (described in Chapter 23) to install GIMP.

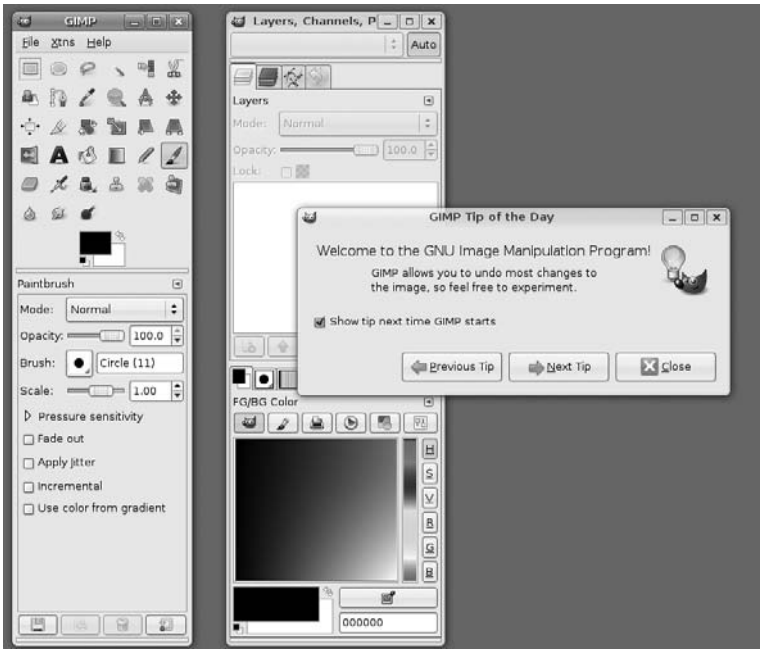
Once GIMP is installed on your system, you can start it manually in a variety of different ways, the most basic of which are:

- By selecting GIMP Image Editor from the Applications ➪ Graphics menu on an Ubuntu system
- By selecting GIMP Image Editor from the K Menu ➪ Graphics menu on a Kubuntu system
- By typing the gimp command in a GNOME Terminal or xterm window and pressing Return

When you start GIMP for the first time, the default display that you see depends on the version of GIMP that you are running. On Ubuntu 8.04 and earlier systems, the main GIMP window and various GIMP panels are displayed, as shown in Figure 18-1.

FIGURE 18-1

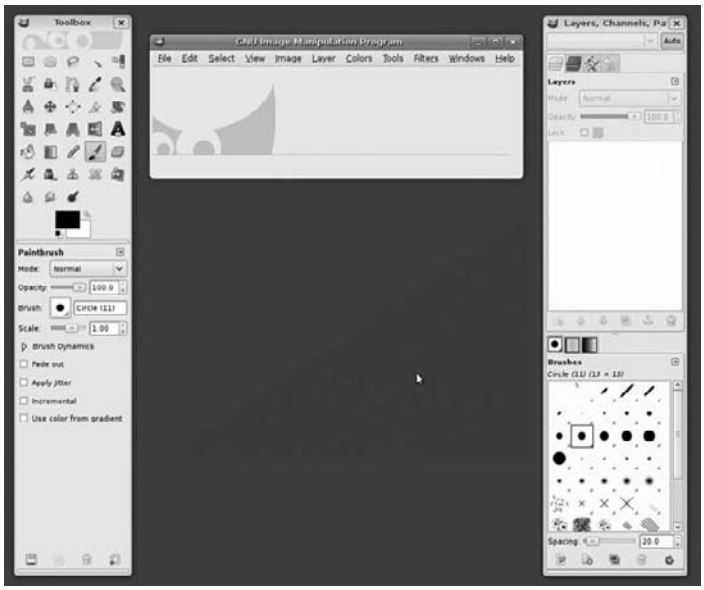
Starting GIMP for the first time on Ubuntu 8.04 or earlier



Ubuntu 8.10 introduces the latest, greatest version of GIMP, GIMP 2.6. This new version provides some basic simplification of the user interface among its many other enhancements. If you are using GIMP on an Ubuntu 8.10, Kubuntu 8.10, or later system, you'll see a startup screen like the one shown in Figure 18-2.

FIGURE 18-2

Starting GIMP for the first time on Ubuntu 8.10 or later



Despite the initial visual differences between Figures 18-1 and 18-2, these versions of GIMP are really very similar. If you're new to GIMP, you can skip to the next section because you probably don't care what it used to look like. If you're already familiar with GIMP, you'll still be able to locate and use your favorite GIMP commands — the primary change in the new version is that a new GIMP window replaces the Toolbox window as the primary GIMP window. The menus that were formerly located on the Toolbox window are now present on the main GIMP window, and have been redesigned to better organize sets of related commands. The main GIMP window serves as your primary image editing window, which was formerly done by separate edit windows. When you create or open a new graphic, it displays in the main GIMP window. Closing the last graphic that you're editing re-displays the generic GIMP window shown in Figure 18-2.

NOTE **If you're a long-time GIMP user, you may remember that GIMP formerly displayed a sequence of startup dialogs that enabled you to customize some aspects of GIMP's initial configuration. Because most people simply accepted all of the defaults, these introductory panels have been done away with. If this is your first time starting GIMP, things are simpler now.**

A Quick Tour of GIMP

The first thing that most people do when using GIMP is to customize the various panels and dialogs that display each time GIMP is launched.

The version of GIMP provided with Ubuntu 8.04 or earlier systems displays a tip of the day when it first starts, and provides the option to automatically see a GIMP tip each time you start GIMP. This default tip display has been removed in GIMP 2.6. If you are using an earlier version of GIMP, you can disable the

startup tips by de-selecting the Show tip next time GIMP starts checkbox. If you want tips at any time in any version of GIMP, you can access them by selecting GIMP's Help ➪ Tip of the Day menu command.

As you can see from Figures 18-1 and 18-2, GIMP initially displays multiple windows. A single window can actually contain multiple dialogs that are combined into one dialog — such dialogs are known as *dockable* dialogs. To make things even more flexible, dialogs can be either docked together or combined as multiple tabs within a single dialog. The default screen shown in Figure 18-2 displays the following windows. Each window description includes a list of the dialogs contained in that window:

- **Main GIMP window:** (GIMP 2.6 and later only) Provides the GIMP menus and shows the first graphic or image that you are editing. Opening another graphic at the same time displays that graphic in a copy of this window.
- **Toolbox window:** The window at the left contains two dialogs: at the top, the Toolbox dialog and below it, the Tool Options dialog. The Toolbox dialog shows the editing or selection tool that is currently selected, and provides access to all other editing and selection tools. The Tool Options dialog provides access to the configurable aspects of the GIMP tool that is currently selected in the Main Toolbox dialog. In versions of GIMP prior to version 2.6, the Toolbox window contains a menu bar that provides access to all GIMP commands and dialogs; this menu bar is located in the main GIMP window in GIMP 2.6 and later.
- **Layers and Brushes window:** The window at the right contains two dialogs: at the top, the Layers, Channels, and Paths dialog and below it, the Brushes, Patterns, and Gradients dialog. The interesting thing about this window is that each of the dialogs that it contains actually consists of multiple dialogs that have been docked together as tabs. The top dialog consists of the Layers, Channels, and Paths and the Undo dialogs docked together, while the bottom dialog consists of the Brushes, Palettes, and Gradients dialogs docked together.

TIP **GIMP's Main Toolbox dialog must always be visible in any version of GIMP — closing this dialog (or the main GIMP window in GIMP 2.6 and later) causes GIMP to exit. Similarly, most GIMP users find it extremely useful to always have the Layers dialog available so that they have easy access to the portions of an image that are contained on different layers.**

Docking and undocking dialogs is a fairly simple process. To undock dialogs, select the header for that dialog (the header inside the window, not the title bar for the window itself), and drag it out of the window in which it is currently displayed. For example, to undock the Toolbox Options dialog from the Toolbox window, left-click and hold on the header of the Toolbox Options dialog (which initially says Paintbrush because the paintbrush is the tool that is initially selected when you first start GIMP), and drag the header out of the Toolbox window. An icon identifying that dialog follows the cursor (which changes into a hand), as shown in Figure 18-3, in which the Tool Options dialog header has been dragged out of the window shown at the left in Figure 18-2.

Releasing this icon on the screen at the point where you want the undocked dialog to appear results in a screen like that shown in Figure 18-4. The Tool Options dialog displays as a separate window, and the portion of the Main GIMP Toolbox window that it formerly occupied is now empty. You can resize the window that formerly contained this dialog to remove the empty space by selecting and dragging a corner.

FIGURE 18-3

Undocking a GIMP dialog

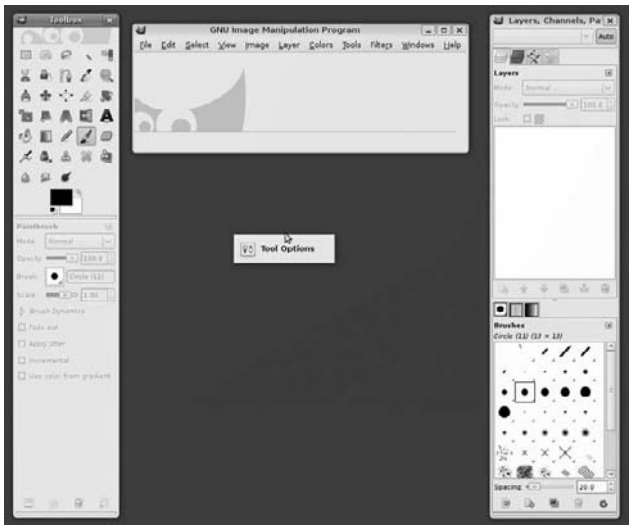
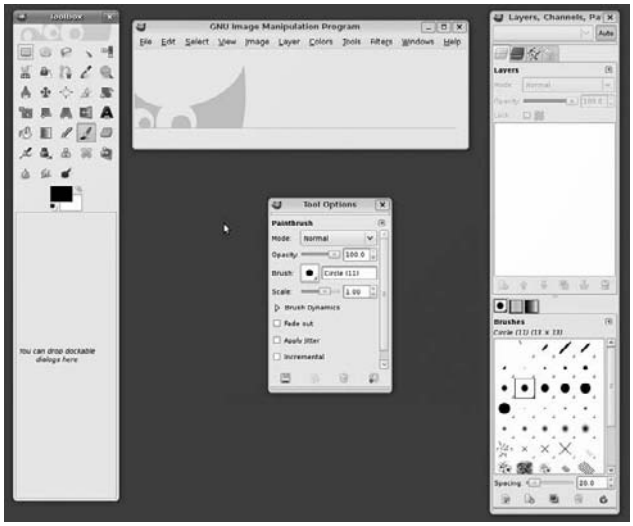


FIGURE 18-4

The GIMP Tool Options dialog, undocked



Color and GIMP

GIMP is often referred to as an Adobe Photoshop killer because it provides much of the functionality of Adobe Photoshop at no cost to the user. For most of us, this is probably true. However, professional illustrators will probably either say "What's GIMP?" or simply laugh when you say "Illustrators of the world unite, and cast off your copies of Photoshop!" Two color-related issues prevent GIMP from serious commercial use: GIMP currently supports only 8-bit color specifications, and GIMP's color handling is fundamentally based on RGB colors, while professional printing relies on CMYK and Pantone color specifications. While work on 16-bit color support in GIMP is in progress (through the new GEGL rendering engine, `www.gegl.org`), it isn't available now, but the lack of built-in CMYK support is the real stumbling block to using GIMP for commercial illustrations that may end up in print publications. Newer versions of GIMP will show both the RGB and CMYK color values for any portion of an image, but GIMP saves only RGB color designations in its output. While most print shops can convert RGB to CMYK for you, this isn't an exact process, typically adding some black to an image when trying to match the colors. Not a good thing for professionals who expect their images to look like their original designs.

Fortunately, GIMP's plug-in architecture lends itself to addressing this sort of problem, and various GIMP plug-ins for CMYK support are available (although they are not installed with GIMP by default). The best known of these for the latest version of GIMP is the Separate+ plug-in by Yoshinori Yamakawa, which is available from `http://cue.yellowmagic.info/softwares/separate.html`. If you are an illustrator, I'd love to hear whether this plug-in can crack the professional printing barrier for GIMP. Crossing fingers now…

Sample GIMP Tasks

As I acknowledged earlier, I am not a graphics artist. Many people who are serious artists have created extensive sets of tutorials on using GIMP, and I'd be insulting them (and you) by either parroting those or by trying to pretend that I have major graphics skills. That said, I do use GIMP regularly for a variety of common tasks that you may find useful, so I'll discuss some of those in the next two sections. For tutorials on using GIMP to do serious graphics work, see some of the URLs in the section "More Information About GIMP," later in this chapter.

Taking Screenshots Using GIMP

As you may have noticed so far, this book contains several screen captures that show the entire screen, specific windows, or certain portions of the screen. My graphics friends have finally convinced me that a picture is indeed worth a thousand words, or should at least accompany them. All of the screenshots in this book were done with GIMP, which is an excellent tool for screen and window captures. You may think that you'll never need to do a screen capture yourself, but that's not necessarily the case — screen captures are very useful when you're submitting problem reports and trying to share what you're seeing on your screen with a friend, and so on.

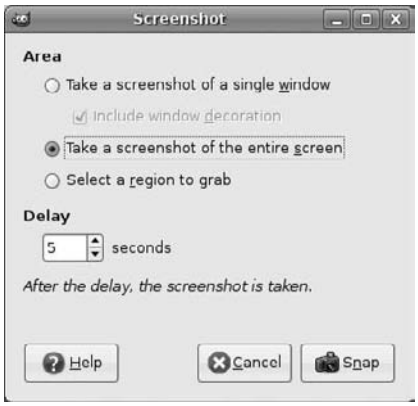
TIP **GIMP isn't the only screen capture software in town — after all, this is Linux. Both Ubuntu and Kubuntu provide lighter-weight screen-capture packages. On Ubuntu systems, you can capture full-screen or single window images by selecting the Applications ➪ Accessories ➪ Take Screenshot menu command (or by running the `gnome-screenshot` command from any terminal window). On Kubuntu systems, you can capture full-screen, single window, or screen region images by selecting the K Menu ➪ Graphics ➪ KSnapshot (Screen Capture Program) menu item (or by running the `ksnapshot` command from any terminal window).**

Capturing a screenshot in GIMP is quite easy:

1. Select File ➪ Create ➪ Screen Shot from the main GIMP window (or the File ➪ Acquire ➪ Screen Shot menu command from the GIMP Toolbox window if you are using an earlier version of GIMP). The dialog shown in Figure 18-5 displays.

FIGURE 18-5

The Screenshot dialog



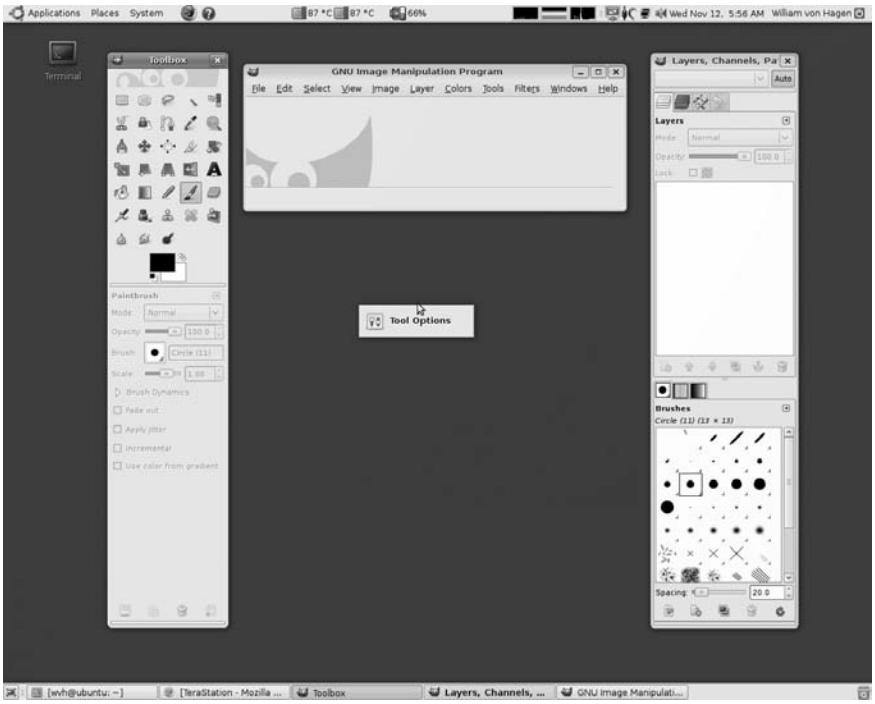
2. Select the radio button that identifies whether you want to capture a specific window or the entire screen.

TIP **GIMP doesn't capture the cursor when taking a screenshot — if you need to do a screenshot that includes the cursor on an Ubuntu system, you can use the gnome-screenshot tool instead. Unfortunately, Kubuntu's equivalent KSnapshot screen capture tool doesn't enable you to capture the cursor.**

3. Select any delay that you want to take place before you capture a specific dialog or the entire screen. Specifying a delay enables you to pull down menus, set specific values, hide other dialogs, and so on — anything that you want to do to capture a screen or window shot that displays exactly what you want it to.
4. Click Snap to initiate the capture process. After any interval that you have specified elapses, the screen capture takes place. (If no delay was specified, it happens as soon as you click Snap.) If you specified that you wanted to capture a single window, the cursor turns into a set of crosshairs, and you must click on the window that you want to capture. Once the capture takes place, the result displays on the screen in its own GIMP window, as shown in Figure 18-6, which shows the entire screen capture for Figure 18-3.

FIGURE 18-6

A full screen capture in GIMP



5. Once you have a full screen or window capture, the next step is to save it in a graphics format that you or someone else can use. To save the screen capture, right-click in the capture window and select File ➪ Save As. (You can also do this by pressing Ctrl+Shift+S, which is certainly my favorite method because of its brevity.) A dialog like the one shown in Figure 18-7 displays.

6. Dialogs like the one shown in Figure 18-7 enable you to save your screenshot in a specific graphics format, which is determined by the file extension. After entering the file name and the file extension that you want to use, click Save to display an additional, format-specific dialog, as shown in Figure 18-8.

TIP **By default, the file extension that you specify when saving a file in GIMP automatically identifies the graphics format in which you want the file to be saved. If you don't know the file extension that is associated with the graphics format in which you want to save your file, click the Select File Type control in the lower-left corner of the Save Image dialog to display a list of available graphics formats. Selecting a graphics format automatically adds the default file extension that is associated with the selected graphics format to the name of the file that you are saving.**

7. The options associated with each graphics format differ, so the available options on dialogs like the one shown in Figure 18-8 are specific to your target file format. Once you have made any changes to the default values that GIMP associates with your target output format, click Save.

FIGURE 18-7

GIMP's first Save Image dialog

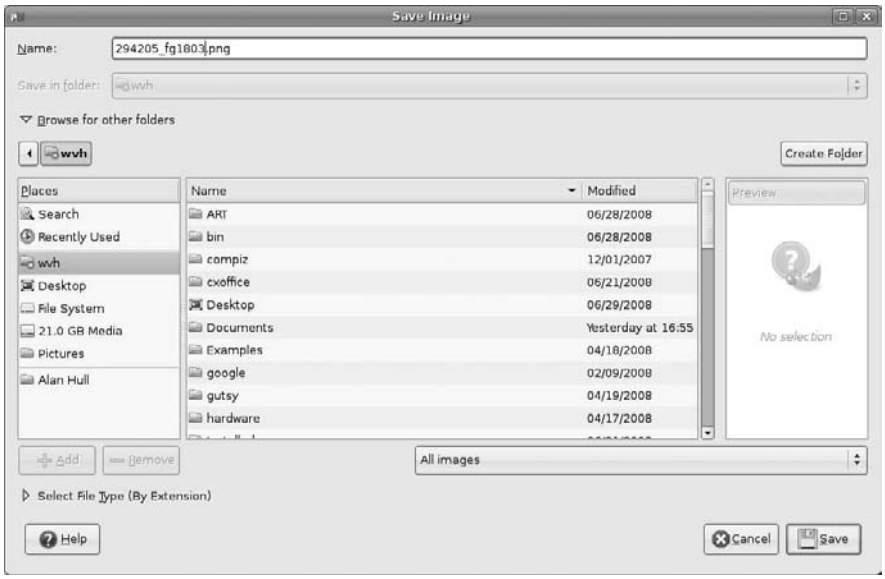


FIGURE 18-8

GIMP's format-specific Save dialog



GIMP not only supports a wide variety of output graphics formats, but it also makes it quite easy to capture high-quality screen images and fine-tune the characteristics of the files that you create.

Converting and Reducing Images Using GIMP

Most of us nowadays have a Web page that we use to share information and pictures with our friends. Most digital cameras save images in JPEG format, which can provide a great balance between small file size (so that Web pages containing images will load in our lifetimes) and good quality (so that you appear to know what you're doing with your digital camera). Regardless, there are going to be times when you have images in larger file sizes that you want to reduce in size and convert to other formats. For example, the TIFF files produced by most scanning software today are huge and can't be displayed on a Web page. Similarly, most digital cameras today capture photos in many megapixels, which are great for printing but are overkill on a Web page.

Because GIMP supports so many different graphics file formats, it is an excellent tool for converting images to formats that you can use elsewhere (most notably on the Web), and for reducing file size without sacrificing too much quality. The sidebar "Graphics Formats Supported by GIMP" shows all of the graphics file formats and associated file extensions that are supported by GIMP.

The following are some tips on using GIMP to reduce the file size of your images so that you can easily put them on a Web page or even mail them to loved ones:

- Save your files in a lossy format that supports compression. PNG files are an attractive choice because they support up to 48-bit color, but I've found JPEG files easiest to exchange with users of other types of computer systems (and they are also supported in all Web browsers — some older browsers don't support PNG graphics). Make sure to save the modified files with a different name so that you still have your originals in case you want to print them or edit them in some other tool — as noted earlier, the original file can never be completely reconstructed from the compressed version.
- When editing or converting JPEG images, experiment with reducing the quality level of your image. In most cases, you won't be able to tell the difference, and neither will your viewers. To reduce the quality level of a JPEG image, click File ➪ Save As, enter the new, unique name, and click Save. The Save As JPEG dialog displays, as shown in Figure 18-9. By default, this dialog selects an image quality level of 85. Click the Show preview in image window checkbox and experiment with dragging the slider to see how your image changes and how the file size changes with lower quality images. Once you're happy with the new image, click Save.
- For Web graphics, reduce the size of your images by scaling them down from the physical size captured by your camera or scanner to a size that is better suited to viewing on a Web page. To scale images in GIMP, select Image ➪ Scale Image. The dialog shown in Figure 18-10 displays. Enter a new value in the Width field and press Return — the Height value will change automatically to preserve the current aspect ratio of your image. Click Scale to rescale the image and update the preview of your image.

Using these simple techniques, I've been able to reduce images from the (approximately) 500KB size that my digital camera captures down to less than 50KB, with no obvious loss of quality. Well, maybe someone with better eyesight than mine could notice a significant difference, but I certainly can't.

TIP **Ubuntu and Kubuntu also provide several command-line tools that can quickly and easily manipulate images. One of my favorites is the `convert` utility, which reduces the color depth of an image. The `convert` utility is one of the graphics tools provided in the ImageMagick set of graphics tools that are installed by default on your Ubuntu system. (You'll have to manually install the `imagemagick` package on Kubuntu systems.) For more information about the convert utility or ImageMagick in general, use the `man` command to see the online reference information for these commands.**

FIGURE 18-9

Reducing JPEG quality when saving files



Graphics Formats Supported by GIMP

GIMP enables you to read and save graphics files in the following formats, each identified by one or more of the following associated file extensions:

- **Adobe Photoshop image:** `.psd`
- **Alias/Wavefront PIX image:** `.pix`, `.mask`, `.matte`, `.alpha`, `.als`
- **ASCII Art:** `.txt`, `.ansi`, `.text`
- **Autodesk FLIC animation:** `.fli`, `.flc`
- **Bzip archive:** `.xcf`, `.bz2`, `.bz2`, `.xcfbz2`
- **C source code:** `.c`
- **C source code header:** `.h`
- **Colored XHTML:** `.xhtml`
- **Digital image and Communications in Medicine image:** `.dcm`, `.dicom`
- **Encapsulated PostScript image:** `.eps`
- **Flexible Image Transport System:** `.fit`, `.fits`
- **GIF (Graphics Interchange Format) image:** `.gif`
- **GIMP Brush:** `.gbr`
- **GIMP Brush (animated):** `.gih`
- **GIMP Compressed XJT image:** `.xjt`, `.xjtgz`, `.xjtbz2`
- **GIMP pattern:** `.pat`
- **GIMP XCF image:** `.xcf`
- **Gzip archive:** `.xcf.gz`, `.gz`, `.xcfgz`

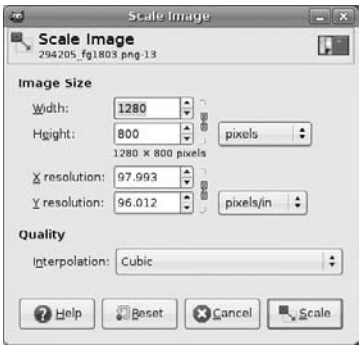
continued

continued

- **HTML table:** `.html`, `.htm`
- **JPEG (Joint Photographic Experts Group) image:** `.jpg`, `.jpeg`, `.jpe`
- **KISS CEL:** `.cel`
- **Microsoft Windows BMP image:** `.bmp`
- **Microsoft Windows icon:** `.ico`
- **MNG animation:** `.mng`
- **PBM (Portable BitMap) image:** `.pbm`
- **PGM (Portable GrayMap) image:** `.pgm`
- **PNG (Portable Network Graphics) image:** `.png`
- **PNM (Portable AnyMap) image:** `.pnm`
- **PostScript document:** `.ps`
- **PPM (Portable PixMap) image:** `.ppm`
- **Silicon Graphics IRIS image:** `.sgi`, `.rgb`, `.bw`, `.icon`
- **Sun Rasterfile image:** `.im1`, `.im8`, `.im24`, `.im32`, and so on
- **SVG (Scalable Vector Graphics) import/read only:** `.svg`
- **TarGA image:** `.tga`
- **TIFF (Tagged Image File Format) image:** `.png`, `.pngf`
- **X Bitmap image:** `.xbm`, `.icon`, `.bitmap`
- **X Pixmap image:** `.xpm`
- **X Window dump:** `.xwd`
- **Zsoft PCX image:** `.pcx`, `.pcc`

FIGURE 18-10

Scaling images for better Web-page display



GIMP has many, many more great features and capabilities — this section has only scratched the surface to encourage you to experiment. See the next section for links to some great online sites that provide detailed information and tutorials about GIMP. Tutorials provide a great way to accomplish a specific task quickly and gain confidence in your growing proficiency.

TIP **GIMP provides a number of default keyboard shortcuts, and also enables you to define your own shortcuts, as described in the "Pimp My GIMP" section of the GIMP documentation at** `http://epierce.freeshell.org/gimp/gimp_ps.php`**. If you're an Adobe Photoshop fan, you may be interested in the set of keyboards shortcuts by E Pierce as a starting point — although these are for older versions of GIMP and Photoshop, they still work in the current version of GIMP, for the most part. Actually, if you're an Adobe Photoshop user, you may be interested in GIMPshop (**`http://plasticbugs.com/?page_id=294`**), which is a version of GIMP in which the menus and dialogs have been reorganized to mimic the Photoshop interface.**

More Information About GIMP

Hundreds of Web sites provide information about using GIMP. The following are some of my favorites:

- `www.gimp.org`: The GIMP home page. Because this is the GIMP mother ship, it's a great starting point for checking the latest documentation, making sure that you have the latest version of GIMP, finding GIMP tutorials, and getting pointers to other online resources. All of the GIMP documentation is available at `docs.gimp.org`. A variety of GIMP tutorials are available at `www.gimp.org/tutorials`.
- `www.gimp.org/links`: Links ranging from the GIMP mothership to interesting GIMP-related sites.
- `http://gimp-savvy.com/BOOK/index.html`: "Grokking the GIMP" is an excellent, although slightly older, book about GIMP that is available online. Although it doesn't have the latest and greatest GIMP information, it is still an excellent resource (and a book well worth having if you're going to be doing lots of work with GIMP).
- `http://gimps.de/en/tutorials/gimp/picture-photo-image/index.htm`: A great set of GIMP tutorials for photo retouching.

Using Krita

Krita (`www.koffice.org/krita`) is a KDE-based image creation and editing tool that many people prefer to GIMP (especially KDE purists who don't want to have to install GIMP's GTK baggage on their shiny KDE boxes). As you may have noticed, I'm not a purist — I believe in installing the best tool for the job on my Linux systems, and that disk space is cheap.

So why use Krita rather than GIMP? If you're a Kubuntu user, Krita uses the KDE libraries and GUI conventions, and works seamlessly with other KOffice applications. More important, graphics professionals on either Kubuntu or Ubuntu systems will especially appreciate its integrated support for the CMYK color specifications that most commercial print shops require. Krita was one of the first open source image applications to provide support for CMYK (which is still somewhat lacking in GIMP, as discussed in the sidebar of the GIMP section of this chapter). Krita also supports 16-bit color (RGB, CMYK, and a few other abbreviations) instead of the 8-bit color support that is currently provided by GIMP. File-format–wise, Krita doesn't support as many import or export formats, but it does support all of the ones that most people care about, including some that GIMP does not support. See the sidebar "Graphics Formats Supported by Krita" for detailed information.

Graphics Formats Supported by Krita

Krita takes a very clever approach to supporting a large number of import and export image formats. Krita itself supports a relatively small number of image formats, and also uses standard KOffice import modules for various formats. To support most other image formats, Krita uses a plug-in to provide access to the huge number of image formats that are supported by an external package known as GraphicsMagick (`www.graphicsmagick.org`). The combination of these enables you to work with graphics files in the following formats, each identified by one or more associated file extensions. This list also identifies whether Krita can both read and write each format.

- **Adobe Illustrator document (import only):** `.ai`
- **Adobe Photoshop image (import only, Photoshop version 6 or earlier):** `.psd`
- **Encapsulated PostScript image (import only):** `.eps`
- **Industrial Light and Magic OpenEXR image:** `.exr`
- **JPEG (Joint Photographic Experts Group) image:** `.jpg`, `.jpeg`, `.jpe`
- **KFormula Document (import only):** `.kfo`
- **KPresenter Document (import only):** `.kpr`
- **Karbon14 Document (import only):** `.kar`
- **Kontour Document (import only):** `.kon`
- **Krita Document:** `.kra`
- **MagicPoint Presentation (import only):** `.mgp`
- **MathML Document (import only):** `.mml`
- **Microsoft Office Drawing (import only):** `.vsd`
- **Microsoft Windows BMP image (import only):** `.bmp`
- **Microsoft Windows icon (import only):** `.ico`
- **Microsoft Windows Metafile (import only):** `.wmf`
- **Microsoft Powerpoint Presentation (import only):** `.ppt`
- **OpenOffice.org Drawing (import only):** `.odg`
- **OpenOffice.org Presentation (import only):** `.odp`
- **OpenOffice.org Presentation Template (import only):** `.odt`
- **PDF (Portable Document Format) Document (import only):** `.pdf`
- **PNG (Portable Network Graphics) image:** `.png`
- **PostScript document (import only):** `.ps`
- **RAW Digital Cameria image (import only):** `.raw`
- **SVG (Scalable Vector Graphics) file (import only):** `.svg`
- **TIFF (Tagged Image File Format) image:** `.tif`, `.tiff`
- **XML Document (import only):** `.xml`

continued

continued

CAUTION **The version of Krita in the Ubuntu repositories at the time of this writing does *not* provide the ImageMagick plug-in, and therefore most attempts to import files in formats other than JPEG, OpenEXR, PDF, PNG, RAW, SVG, and TIFF will fail. Quite bogus! If this hasn't been fixed by the time that you read this, you can get around the problem by installing the GraphicsMagick package and using its command-line gm convert utility in a terminal window to convert files from other formats into a supported format. For a list of the import and export formats that are supported by GraphicsMagick, see** `www.graphicsmagick.org/www/formats.html`.

Starting Krita

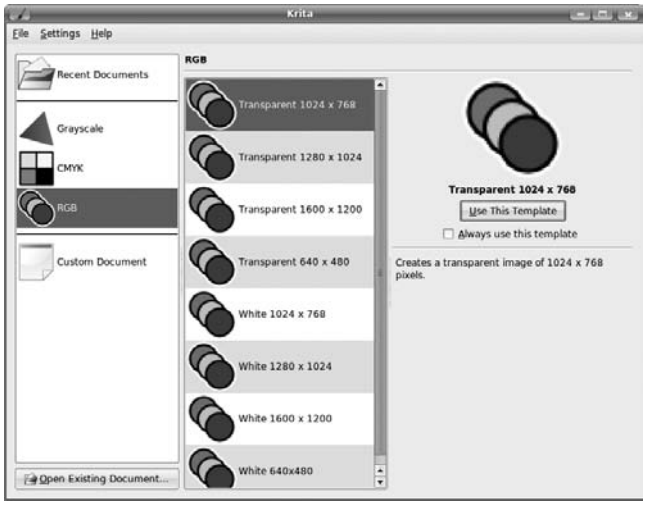
Krita is a component of the KOffice suite, but is not installed by default on Kubuntu systems and is obviously not installed by default on GNOME-based Ubuntu systems. You can install Krita on a Kubuntu or Ubuntu system using any of the tools discussed in Chapter 23. When installing Krita on a system, you may also want to install the krita-plugins package, which is not automatically installed with Krita. As the name suggests, this package provides some handy plug-ins for Krita, including ones for dodging and burning portions of your images (i.e., lightening and darkening portions of the images), image blurring, grayscale conversion, and red eye removal.

To start Krita on a Kubuntu system, select K Menu ➪ Graphics ➪ Krita Painting and Image Editing. To start Krita on an Ubuntu system, select the Applications ➪ Graphics ➪ Krita menu item. You can also start Krita by typing the krita command in a GNOME Terminal, Konsole, or xterm window and pressing Return.

When Krita starts, it displays its image wizard, shown in Figure 18-11.

FIGURE 18-11

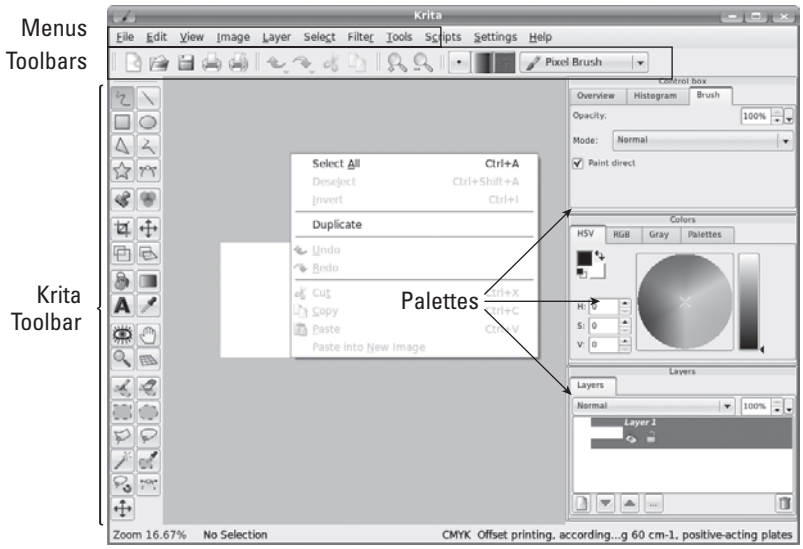
Krita's image wizard



Krita's image wizard enables you to select from different image templates for any new image that you are creating, and also provides the Recent Documents button and Open Existing Document dialog to make it easy to re-open a Krita image that you've worked on before. Krita provides single grayscale (640 × 480) and CMYK (2000 × 800) templates, as well as a variety of transparent and white-background RGB templates ranging in size from 640 × 480 to 1600 × 1200. If you are creating a new image, select the type of (and size of, if necessary) template that you want to use and click Use This Template. Krita starts and displays a window like the one shown in Figure 18-12. This particular window uses the CMYK template.

FIGURE 18-12

An initial Krita screen



A Quick Tour of Krita

Like all KOffice tools, Krita provides excellent online help, so I won't bore you by walking through each menu and entry that it contains. However, as a quick overview of the Krita interface, the callouts in Figure 18-12 highlight the following areas:

- **Menus:** The drop-down menus in Krita are organized much like those in other drawing packages to provide a familiar user experience. Menu items with an arrow at their right edge lead to other, related submenus.
- **Toolbars:** By default, Krita displays the File, Edit, Navigation, and Brushes and Stuff toolbars side-by-side below the menu bar. Krita's Krita toolbar is shown to the left of the image editing area. You can modify the list of icons displayed on a toolbar by right-clicking to the right of these toolbars and selecting the Configure Toolbars command from the context-sensitive menu that displays, or by selecting the Configure Toolbars command from the Settings ➪ Toolbars menu. You can then select the toolbar that you want to edit from the drop-down list at the top of the dialog that displays. You can deactivate or activate toolbars by selecting Settings ➪ Toolbars,

and clicking the name of the toolbar that you want to deactivate or activate from the menu that displays.

- **Palettes:** Krita displays specialized tool dialogs known as palettes to the right of the image editing area, each of which can contain multiple tabs. These palettes are as follows:
 - **Control Box:** By default, Krita's Control Box palette displays a Brush tab because the Brush tool is selected when you first start Krita. Selecting other tools from the Krita toolbar to the left of the image editing area displays tabs that are relevant to the selected tool. The Control Box can also display an Overview tab that provides a thumbnail of the entire image (which can be useful when you have zoomed in on an image) and a Histogram tab that provides a histogram of color use in the current image.
 - **Colors:** By default, the Colors palette displays HSV, RGB, Gray, and Palettes tabs to enable you to set the current fill and drawing colors. The Colors palette can also display a Watercolors tab that enables you to simulate painting with wet paint. (This is a great effect that is well worth experimenting with!)
 - **Layers:** By default, the Layers palette displays a Layers tab that contains an entry for each layer that is in use in your document, each of which enables you to show, hide, and lock that layer by clicking on the icons in the entry for that layer. The Layers palette can also display a Script Manager tab, which enables you to select and execute image manipulation scripts written in Python or Ruby.

 You can show or hide the Palettes window or the tabs in different palettes by selecting or deselecting them using the commands on the View ➪ Palettes submenu.
- **Status Bar:** Shown at the bottom of the Krita window, the status bar provides information about the current Zoom, Selection, and Colorspace settings.
- **Context-sensitive menu:** Right-clicking on the background of any image that you are working on in Krita displays the context-sensitive menu shown in Figure 18-12. This menu provides quick access to several commonly used commands and submenus, including those for selecting and modifying the current layer or a portion of an image, as well as providing access to standard cut and paste commands.

Krita provides extensive online help, available by selecting Krita Handbook from the Help menu, or by pressing the F1 key on your keyboard at any time. Figure 18-13 shows the main page of the Krita handbook.

Test-Drive: Using Krita

Discussing every nuance of creating and editing images in Krita is beyond the scope of this book, and probably deserves its own book. However, to give you a feel for using Krita, this section provides an example of using Krita to import an SVG image (which will use the RGB colorspace by default), changing the image to use the CMYK colorspace, and then saving it as a TIFF file that would be suitable for sending to a commercial printer.

NOTE **Although the grayscale screenshots in this book won't make it very obvious, this example highlights the problems encountered when simply converting RGB to CMYK, because changing the colorspace changes the shade of blue used in the example image. If you want to follow along, see the image used in this section, which is the file warszawianka_Krakow_-_coat_of_arms.svg (the coat of arms of the city of Krakow, Poland), which is from the Open Clip Art library. The Open Clip Art Library was discussed in a tip at the beginning of this chapter.**

FIGURE 18-13

Krita's online help

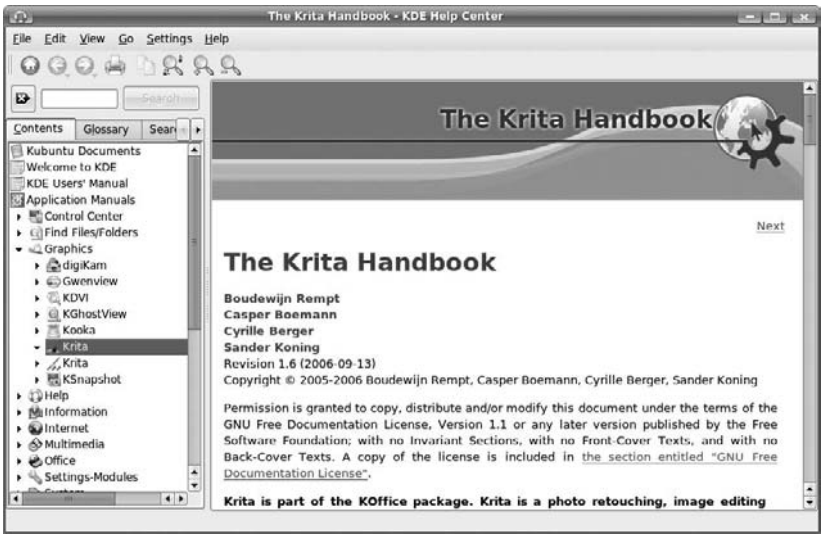
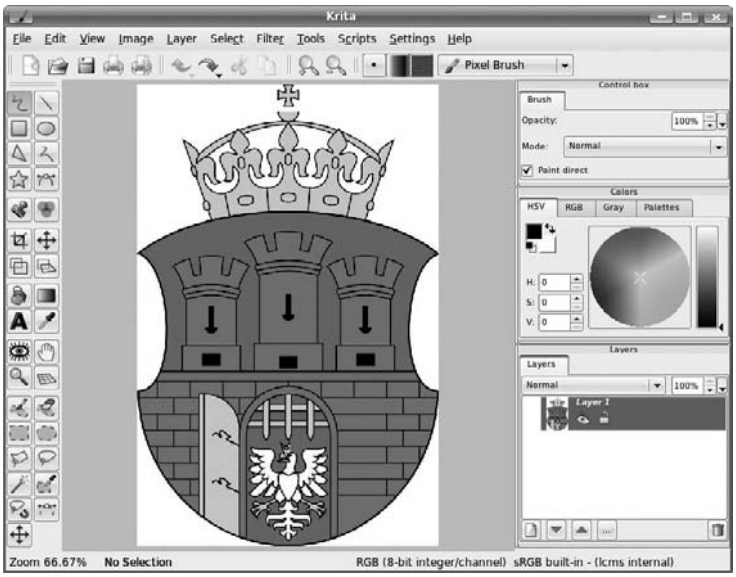


FIGURE 18-14

An imported SVG file shown in Krita



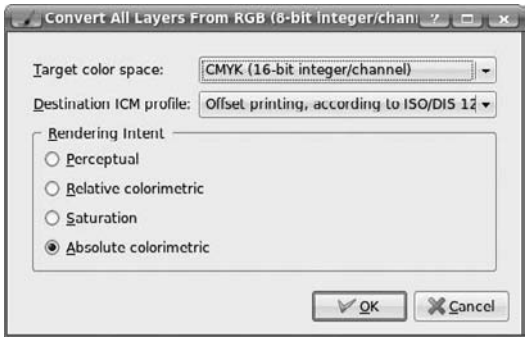
1. Start Krita using one of the methods discussed in the previous section. When the image wizard displays, click Open Existing Document and navigate to and select the file that you want to open. Krita imports and displays the selected file, as shown in Figure 18-14.

 As you can see in the status bar at the bottom of Figure 18-14, the imported file uses the 8-bit RGB colorspace by default.

2. Select the Image ➪ Convert Image Type menu item. The dialog shown in Figure 18-15 displays.

FIGURE 18-15

The dialog for colorspace conversion



3. If you do not want to use the default, which is to convert to 8-bit CMYK, click the Target Color Space drop-down and select the target colorspace that you want to use. In this example, I'll select 16-bit CMYK. This provides a much larger range of control over the colors in your images, helping ensure that the image from the print shop closely matches what you see on the screen.

4. You can modify the Rendering Intent by selecting any of the radio buttons in this portion of the dialog. The rendering intent specifies how colors in the color space that is currently being used in your image will be converted/mapped to those in the color space that you are changing to:

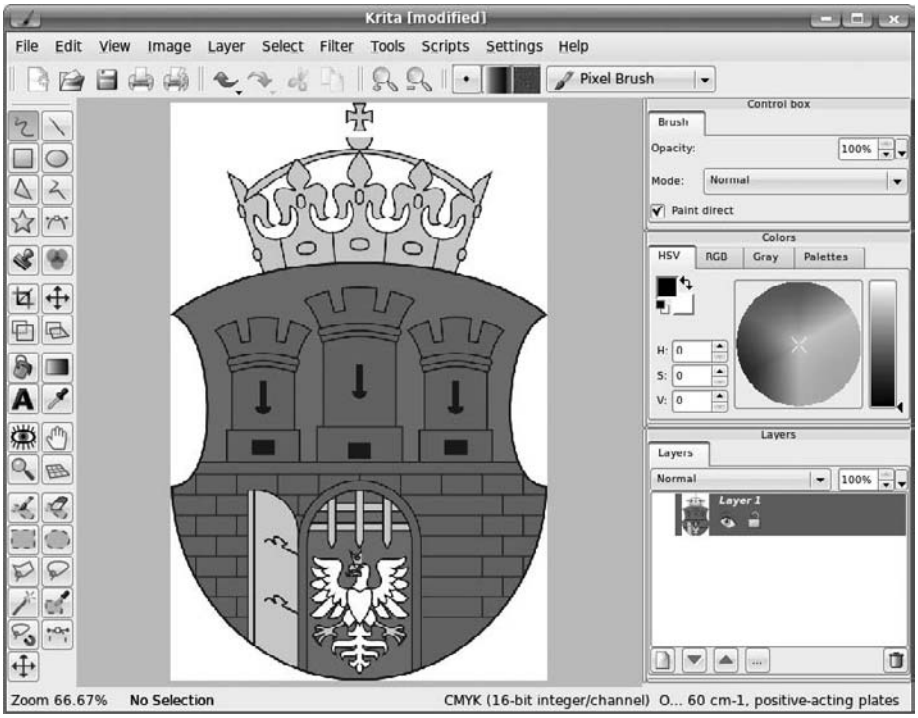
 - **Perceptual:** The existing color space is directly mapped to the target color space. If the two color spaces have different ranges, colors may be shifted when mapping one range to the other.
 - **Relative colorimetric:** Converts each color in the existing color space to the closest color in the target color space. Depending on the range of colors available in each space, this can result in multiple colors being mapped to the same color in the target color space.
 - **Saturation:** Fully saturated colors will remain fully saturated, even if this requires that an existing color be changed.
 - **Absolute colorimetric:** Converts each color in the existing color space to the closest color in the target color space but does not change the value that designates the color white. This may result in unwanted changes to colors, and can also result in multiple colors being mapped to the same color in the target color space.

 In this example, I'll use Absolute colormetric rendering intent because I am going from 8-bit RGB to 16-bit CMYK, and that setting should minimize the amount of color change in the image. Click OK to proceed.

5. The updated image displays, as shown in Figure 18-16.

FIGURE 18-16

The converted image

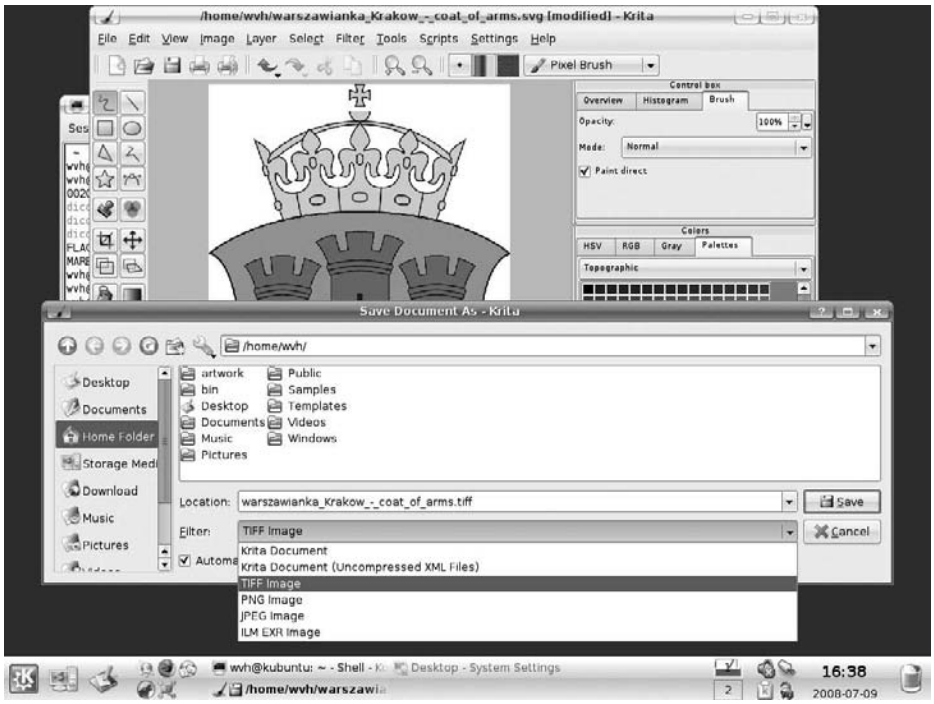


As you can see in the status bar at the bottom of the screen, the image now uses a 16-bit CMYK color space. If you are following along at home (or are really, really good at comparing grayscale screenshots), you will also note that the colors in your image have changed. In real life, you would want to correct the fill and spot colors used in the image, which you could do by using any of the tabs in the Color palette to select a new color, selecting the fill tool (the traditional paint bucket) in the Krita toolbar, and selecting the region that you want to correct.

6. Save the image in an output format that supports CMYK. Select the File menu's Save As command. To specify the new output format, you can use a well-known extension or select the format that that extension represents from the Filter drop-down. I prefer to do the latter because I'm not always sure which extensions are associated with which graphics formats. Click the Filter drop-down and select TIFF, as shown in Figure 18-17. Note that the Krita appends the appropriate file extension to the filename once you specify an output format.
7. Click Save to save your 16-bit CMYK image in TIFF format, and exit from Krita by selecting the File ➪ Quit menu command.

FIGURE 18-17

Specifying the new output format



This simple example is truly the tip of the iceberg. Krita is a powerful application that offers some features that more common applications, such as GIMP, do not. Krita is still maturing; the most recent versions have brought both substantial performance improvements and substantial reductions in memory and system requirements — an unusual (but great!) combination. If you are interested in producing graphics on a Kubuntu or Ubuntu system, and need to be able to give those graphics to a professional print shop, investing a bit of time in using and experimenting with Krita may provide the graphics software solution that you've been searching for. And you can't beat the price!

More Information About Krita

The KOffice suite has been popular for quite a while on KDE-based systems. Krita has been under development since 1999, and a growing number of online resources are available to provide information about using it. The following are my favorite Krita resources on the Web:

- `www.koffice.org/krita`: Krita's home page in the KOffice Web site. This is a good starting point for information about Krita, its interaction with other Koffice applications, the latest version, the Krita FAQ, troubleshooting tips, and much more.

- `http://docs.kde.org/development/en/koffice/krita`: The main documentation site for the Krita Handbook, which is slightly out of date at the time of this writing. Aside from explaining how to use Krita, the handbook also provides some useful tutorials.
- `www.krita-plugins.org`: The home page for the Krita Plugins projects, which is a good spot to watch if you want to make sure that you have the latest and greatest plug-ins and related extensions.

As Krita matures and improves, it has been getting more and more attention in the Linux and graphics press. A quick Clusty or Google search will help you find the latest and greatest information about Krita, hot off the virtual presses.

Photo Editing Overview

Digital photography rules. Although many of us probably still have envelopes stuffed with old holiday and vacation photographs that have been glued together by time and humidity, it's been a long time since I've had to deal with film. I'm probably showing my age by mentioning it at all. I have a few generations of digital cameras with increasing resolution, my cell phone has a camera, my laptop has a camera, and even a wristwatch that I bought during an eBay seizure a few years ago has one. Digital photography makes it incredibly easy to archive, catalog, and share your photos while also reducing the number of items that you can't find when you want them — they're all on your computer. (It also increases the number of items that you should back up!) Having all of your photos as digital images also makes it incredibly easy to "fix" problems ranging from red pupils and various forms of under- or overexposure, to Bolshevik-like purges of estranged relatives or insignificant others from the visual history of your life.

Although I use GIMP for editing all of my digital photographs, both Ubuntu and Kubuntu provide sophisticated software packages that make it easy for you to download images from a digital camera, store them as online albums, and clean up those images in zillions of ways. Google also offers Picasa, a free download that provides a nice solution for retrieving, archiving, and editing photos. (Picasa also runs on Microsoft Windows, in case that still matters to you.) These are discussed in more detail in the section of Chapter 21 on "Digital Cameras and Ubuntu," but I'll at least introduce them here. Ubuntu desktop systems install the F-Spot Photo Manager, shown in Figure 18-18, which is invoked by default when you attach a digital camera to your Ubuntu system, automatically extracts and archives photos, and provides some easy-to-use tools for cropping, removing red-eye, correcting colors, and much more.

CAUTION **F-Spot requires the odious Mono runtime environment that provides support for Microsoft .NET applications on Linux. Perpetuating .NET applications through Mono does not benefit free software. However, if you don't care about that type of issue, F-Spot is a great piece of software. While I may disagree with its developers' choice of programming languages, it's impossible to argue against the quality of the tool.**

Figure 18-19 shows digiKam, the digital photograph editor and archiver that is installed by default on Kubuntu systems. digiKam is an extremely impressive package (with no .NET baggage!) that KDE offers to invoke when you attach a digital camera to your Kubuntu system.

digiKam is a great piece of software, providing a huge number of built-in tools for editing your digital photos, including tools for red-eye correction; color, hue, and gamma correction; resizing and cropping your photos; changing their aspect ratio and perspective; and much more. If you want a dedicated tool for editing digital photos, you can't go wrong with digiKam. Kubuntu users win by default. For Ubuntu users, digiKam is one of the best arguments for installing KDE apps on Ubuntu that I've seen since K3b (discussed in Chapter 19, "Working with Multimedia").

FIGURE 18-18

Ubuntu's F-Spot photo manager



FIGURE 18-19

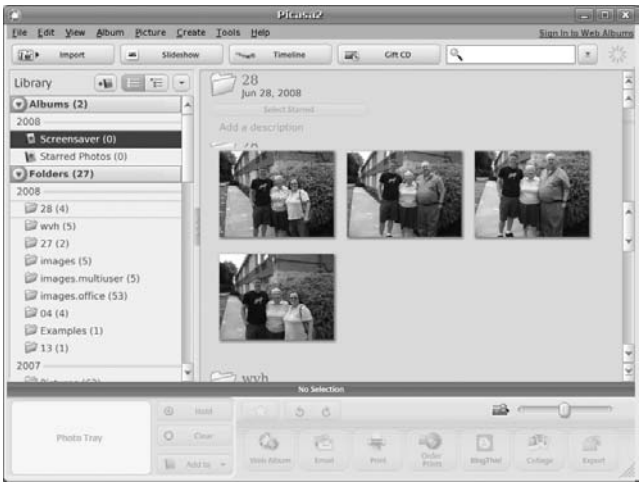
Kubuntu's digiKam photo management software



Everybody knows Google, and most people know that Google provides lots of free desktop and Web-based applications. Unfortunately, many of its desktop applications initially appear for Windows and are only later upgraded to run on Linux (which is somewhat ironic because Goggle is a huge Linux advocate and uses a homebrew Linux distribution on all of its servers). Picasa, Google's photo editing app, uses the software from the WINE project to run on Linux, which is somewhat disappointing but better than nothing. (WINE stands for "WINE is not an Emulator," and is an implementation of the Microsoft Windows API that runs on non-Windows systems.) If you like Picasa, you can also sign up for a Google Picasaweb account (`www.picasaweb.google.com`) to share your Picasa photo albums online. Like the other photo-editing applications introduced in this section, Picasa is discussed in more detail in Chapter 21. Figure 18-20 shows Picasa running on an Ubuntu desktop system.

FIGURE 18-20

Google's Picasa photo management software



Applications such as GIMP and Krita, discussed earlier in this chapter, can largely eliminate most opportunities for Photoshop-envy if you're moving to Linux from some other platform. However, if you simply must use Photoshop, you can always run it under WINE on Ubuntu or Kubuntu desktop systems or run it inside a virtual machine. And there's more! Never one to leave a bandwagon un-jumped-upon, Adobe offers a version of Photoshop as an online photo-editing tool known as Adobe Photoshop Express (`https://www.photoshop.com/express/landing.html`). Signing up for Photoshop Express also gives you 2GB of online storage for your images. On the down side, Photoshop Express requires Flash 9 and I generally dislike Flash on Web sites in much the same way that I dislike cotton candy on the dinner plate. I'm also perfectly happy with the GIMP and Digikam, and am wary of attempts to ensnare my digital content in online tools. Your mileage may vary.

Commercial photo-editing software is also available for Linux — hooray! One up-and-coming photo editing package for Linux is Light Craft's LightZone (`www.lightcrafts.com/linux`), which is a commercial package with a free trial edition. Light Craft seems quite serious about providing high-quality photo-editing software for Linux, and the LightZone software has gotten good reviews and a warm reception from the Linux press. Similarly, an excellent commercial photo-editing package that's well worth

checking out is Bibble (www.bibblelabs.com). Bibble comes in two versions, Lite and Professional, and also offers a trial version. Bibble and LightZone are not included in your Ubuntu Linux installation because they are commercial applications, although both are still relatively inexpensive. Bibble and LightZone are great examples of the growing amount of commercial software for Linux that makes it the equal of any other personal computer operating system.

Using OpenOffice.org Draw

As you may have heard before, OpenOffice.org (www.openoffice.org) is a suite of office applications including a word processor (Writer), a spreadsheet (Calc), a drawing program (Draw), software for creating presentations (Impress), and a database program called BASE. This section focuses on the drawing application — other chapters in this book discuss the OpenOffice.org word processor, spreadsheet, and presentation software. For some background about the OpenOffice.org suite in general, see the introductory section of "Word Processing with OpenOffice.org Writer" in Chapter 16.

Draw is installed by default on desktop Ubuntu and Kubuntu systems. Draw is primarily focused on creating drawings in OpenDocument Drawing (.odg) format. Like Inkscape and Krita (discussed earlier in this chapter), OpenOffice.org Draw enables you to create drawings in SVG format, although its ability to import existing SVG graphics is limited at the time of this writing. Stay tuned!

To avoid repeating the phrase "OpenOffice.org Draw" more times than necessary, I'll refer to it simply as Draw or as oodraw (the name of the application that you must type to execute OpenOffice.org Draw from a command-line interface) throughout the rest of this section.

Starting Draw

To start Draw on an Ubuntu system, select the Applications ➪ Graphics ➪ OpenOffice.org Drawing menu item. To start Draw on a Kubuntu system, select the K Menu ➪ Office ➪ OpenOffice.org Drawing menu command. You can also start Draw by typing the oodraw command in a GNOME Terminal, Konsole, or xterm window and pressing return.

When Draw starts, it displays an empty document, as shown in Figure 18-21.

A Quick Tour of Draw

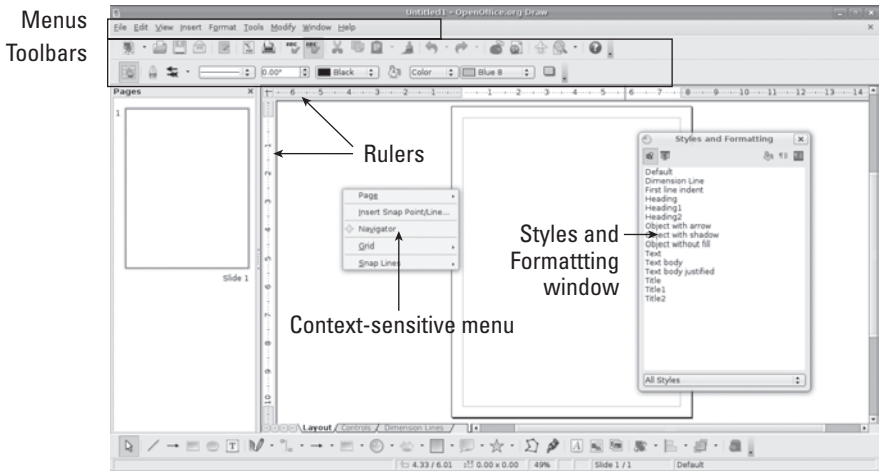
Draw provides excellent online help, as do all of the tools in the OpenOffice suite, so I won't bore you by walking through each menu and entry that it contains. However, as a quick overview of the Draw interface, the callouts in Figure 18-21 highlight the following areas:

- **Menus:** The drop-down menus in Draw are organized much like those in other drawing packages to provide a familiar user experience. Items on submenus with an arrow at their right edge lead to other, related submenus.
- **Toolbars:** By default, Draw displays the Standard, Line and Filling, and Drawing toolbars, which provide one-click access to specific commands when you click on the appropriate icon. You can modify the list of icons displayed on a toolbar by clicking on the control for that toolbar, located at the far right. You can display other toolbars or deactivate the default toolbar by selecting toolbars from the View ➪ Toolbars submenu. Draw provides a great selection of drawing-oriented toolbars, such as a 3D objects toolbar, so that you can quickly and easily create and modify the geometry of 3D images.

- **Rulers:** The rulers display the horizontal and vertical location of the cursor in the current drawing. The grayed-out portions of each ruler reflect the area outside the page margins used in the current document. You can hide these rulers by deselecting the View ➪ Ruler menu command.
- **Styles and Formatting window:** This window is not displayed by default, but you can display it at any time by pressing the F11 key on your keyboard. (Press F11 again to close this window.) The Styles and Formatting window provides easy access to the styles that are available in the current drawing, and is especially convenient if you choose not to display the formatting toolbar. As with any floating window in OpenOffice, you can dock this window by double-clicking on the gray portion of the window heading, making it a part of the default OpenOffice interface rather than a floating window. You can undock a window by repeating this process in the docked area.
- **Context-sensitive menu:** Right-clicking on the background of any drawing that you are working on in OpenOffice.org Draw displays the context-sensitive menu shown in Figure 18-21. This menu provides quick access to several commonly used commands and submenus, including those for modifying the current page format, displaying a grid, and controlling snapping to grid points.

FIGURE 18-21

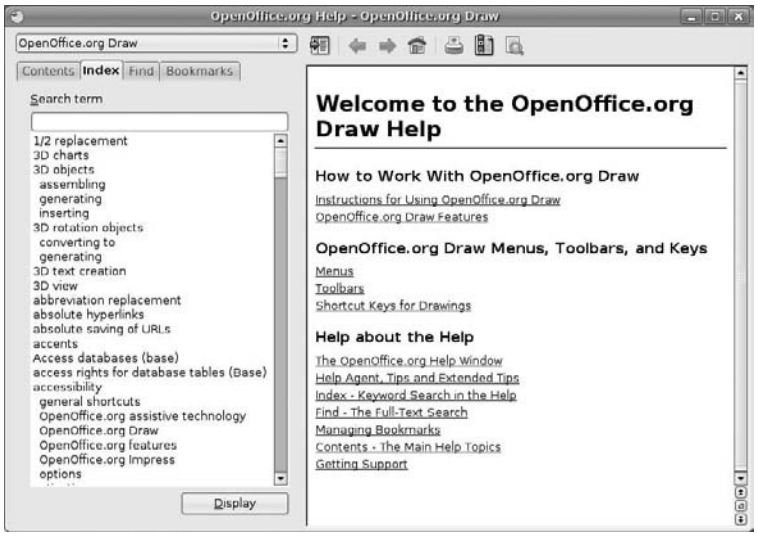
Starting OpenOffice.org Draw



Draw provides extensive online help, available by selecting OpenOffice.org Help from the Help menu, or by pressing the F1 key on your keyboard at any time. Figure 18-22 shows the main page for Draw's help, which provides both a searchable index and a list of popular, general topics from its initial display window (at the right).

FIGURE 18-22

OpenOffice.org Draw's online help



More Information About OpenOffice.org Draw

As the OpenOffice Suite continues to grow in popularity, more and more free resources are available for helping you with each of its components. The following are my favorite Draw-related sites on the Web:

- `http://oooauthors.org/en/members/tutorials/draw/`: Drawing tutorials and Frequently Asked Question (FAQ) files in a variety of languages.
- `http://openoffice.blogs.com/openoffice/draw/index.html`: A blog by Solveig Haugland, author of the OpenOffice.org Guidebook. This portion of the blog contains a few tutorials, a variety of tips and tricks, and pointers to related information such as sources of clip art, templates, and more.
- `www.openoffice.org`: The OpenOffice mother ship, this site provides access to mailing lists, sample documents, FAQs, and much more.

Vector Graphics Tools for Linux

Standard document and graphics formats were less important before personal computers became the latest household appliance and the Internet provided a common platform for exchanging and presenting text and illustrations. Previous chapters stressed the importance and value of open standards for the written word such as HTML and the Open Document format. Similar standards are emerging in the area of graphics and illustration, where older formats such as GIF and JPEG are being complemented by the emergence of open raster graphics standards such as Portable Network Graphics (PNG) and an equally powerful, open standard for vector graphics such as the W3C standard Scalable Vector graphics (SVG) format. (For more information about the SVG standard, see `www.w3.org/Graphics/SVG/`.) Aside

from being designed from the ground up as an open format, one of the primary advantages of SVG files is that they are, as the name suggests, scalable. Because SVG files consist of drawing and color information rather than actual images, they can be resized with no loss in image quality, as you would find in raster graphics formats such as PNG, JPEG, GIF, and so on.

TIP **Although OpenOffice.org Draw, discussed previously, can export images in SVG format and has some import capabilities, Inkscape and Karbon14 were designed for SVG graphics. The OpenDocument Drawing format used by Draw is an excellent format to use if your focus is on producing drawings for use with the other components of the OpenOffice.org suite, in documents, presentations, and so on. If your focus is on creating standalone artwork for Web pages, I'd suggest that you use a program that was designed for vector graphics, such as Inkscape or Karbon.**

SVG graphics files are saved in XML format, which makes them highly portable and very flexible. As mentioned in the section of this chapter on GIMP, most professional printers require CMYK or Pantone color specifications. Both Inkscape and Karbon provide some amount of support for CMYK color specifications, but integrated Pantone support is a legal swamp that no one wants to wade through. However, SVG's fill and stroke color attributes are natural places to attach tags that identify Pantone numbers. If you're going to try this, you should talk this over with your printer (the human, not the device) before simply assuming that the printer knows about and will be able to use this information.

Using Inkscape

Inkscape is an open source application for creating and editing vector graphics, and is similar to traditional applications such as Adobe Illustrator, Aldus (Adobe) Freehand, and CorelDraw. Inkscape can also import graphics in a variety of raster graphics formats that it can use as components of its SVG output. Inkscape can even export graphics in PNG format for compatibility with raster-only applications and older Web browsers. In addition to the standard elements of scalable vector graphics files, Inkscape also supports Web-oriented Creative Commons metadata and direct XML creation and editing.

TIP **When installing Inkscape, you may want to consider installing recommended programs such as dia (a diagramming tool) and the Perl and Python XML extensions at the same time. These provide many import and export extensions for Inkscape, which you'll be glad to have if you ever need them. You can use aptitude's `--with-recommends` option, or select recommended tools from the Synaptic context-sensitive menu. If you're installing Inkscape on a Kubuntu system using the Adept Installer, you can use the Show Details context-sensitive menu command to identify the other packages that you may want to install.**

Because not everyone needs to create vector graphics images (and many people may have already decided to do so using OpenOffice), Inkscape isn't installed by default on Ubuntu or Kubuntu Linux systems. Luckily, that's easy enough to fix, using the standard software installation tools discussed in Chapter 23.

If you're a Kubuntu user and/or a KDE purist, you may want to check out Karbon14, the KDE-based vector graphics program that's discussed in the next section, before installing Inkscape. I use Inkscape for a variety of Web and writing-related tasks, and think it's great, but Linux is all about choice.

Starting Inkscape

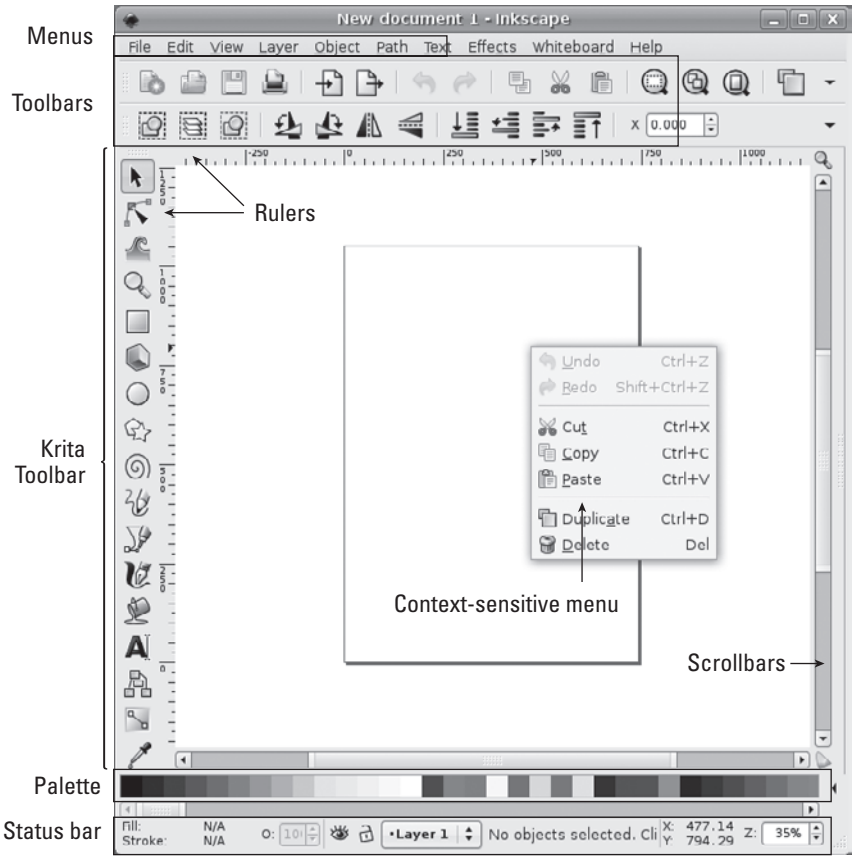
After installing Inkscape, you can start the program in a variety of different ways, the most basic of which are:

- On Ubuntu systems, by selecting Inkscape Vector Illustrator from the Applications ➪ Graphics menu
- On Kubuntu systems, by selecting Inkscape Vector Illustrators from the K Menu ➪ Graphics menu
- By typing the inkscape command in a GNOME Terminal or xterm window and pressing Return

The application starts, displaying an empty document, as shown in Figure 18-23.

FIGURE 18-23

Starting Inkscape



A Quick Tour of Inkscape

Inkscape provides great online help, and even includes some tutorials that are available from its help menu. The callouts in Figure 18-23 highlight the basic components of the Inkscape interface, which are as follows:

- **Menus:** The drop-down menus in Inkscape are organized to provide easy access to commands related to the basic elements of SVG drawings, such as Layer, Object, Path, and Text. The Inkscape designers have done a great job at minimizing the number of submenus used in the program — almost every command is directly available from a top-level menu.
- **Toolbars:** Inkscape provides two horizontal toolbars. The top toolbar, the Commands bar, provides access to some standard file and clipboard manipulation commands, and also provides access to a variety of commands for zooming, copying, and grouping selected objects. The bottom toolbar, the Tool Controls bar, provides easy access to rotation, layer manipulation, and scaling

commands. Clicking the down arrow at the right of each of these toolbars displays additional, related commands for those toolbars.

- **Toolbox:** A vertical toolbox along the left side of the edit window. The toolbox contains icons for common creation and editing commands, such as those for selecting objects; zooming in and out; and creating squares, circles, 3D boxes, various sorts of lines, and so on.
- **Rulers:** These rulers display the horizontal and vertical location of the cursor in the current document.
- **Scrollbars:** Horizontal and vertical scrollbars shown at the bottom and right of an Inkscape window make it easy to move around in zoomed or otherwise large images. If the entire image can be viewed at once, the scrollbars do not appear.
- **Palette:** A horizontal palette at the bottom of an Inkscape window makes it easy to select new drawing and fill colors.
- **Status bar:** The status bar at the very bottom of an Inkscape window provides information about the current layer, selected fill color and stroke type, the current X and Y cursor position, the current zoom, and so on.
- **Context-sensitive menu:** Right-clicking on any component of a drawing that you are working on in Inkscape displays the context-sensitive menu that is also shown in Figure 18-23. This menu provides quick access to several commonly used commands and submenus, including those for duplicating a selected object, grouping and ungrouping, undo operations, and a general object properties panel for the selected object.

You can hide (or subsequently reveal) all of these (except for the menus) by deselecting (or selecting) the associated items from Inkscape's View ➪ Show/Hide menu. You can also configure various Inkscape preferences by selecting the File ➪ Inkscape Preferences menu command. These preferences can be very handy for controlling cursor, scrolling, and selection granularity, which can be critical when doing detailed drawings in Inkscape.

More Information About Inkscape

I'm not the best graphics artist, so I'll leave the evangelizing to those who are. The following are some of my favorite Inkscape-related sites on the Web:

- `www.inkscape.org`: The home page of the Inkscape project, this page provides access to many Inkscape resources, such as mailing lists, articles, collections of images, and so on.
- `www.inkscape.org/mailing_lists.php`: Inkscape-specific mailing lists. The Inkscape-user list is a great place for asking questions and searching for solutions to whatever problem you're experiencing (with Inkscape, that is). The Inkscape-announce list is a low-traffic list, but is worth subscribing to so that you are notified of new major and patch releases of Inkscape.
- `http://tavmjong.free.fr/INKSCAPE/MANUAL/html/index.php`: Tavmjong Bah is writing a book about Inkscape that is freely available on the Web. This is a great book with a fantastic selection of information.
- `www.angelfire.com/mi/kevincharles/inkscape/index.html`: An unofficial (but extremely useful!) user manual for Inkscape written by Kevin Wixson.

As with every popular open source package, more resources and sources of related information for Inkscape are probably appearing on the Web as I type this. If you find a great Inkscape resource that I've missed in this quick summary, please let me know. Inkscape is a great program!

Using Karbon14

Karbon14, also known to its friends simply as Karbon, is a KDE-based application that is a component of the KOffice suite. Like Inkscape, Karbon14 is a vector-based graphics editing application — although it has its own `.karbon` file format, Karbon14 also does a great job of creating and editing graphics in the standard SVG format.

Although it is a component of the KOffice suite, Karbon14 is not installed by default on Kubuntu systems, and is obviously not installed by default on GNOME-based Ubuntu systems. You can install Karbon14 on a Kubuntu or Ubuntu system using any of the tools discussed in Chapter 23.

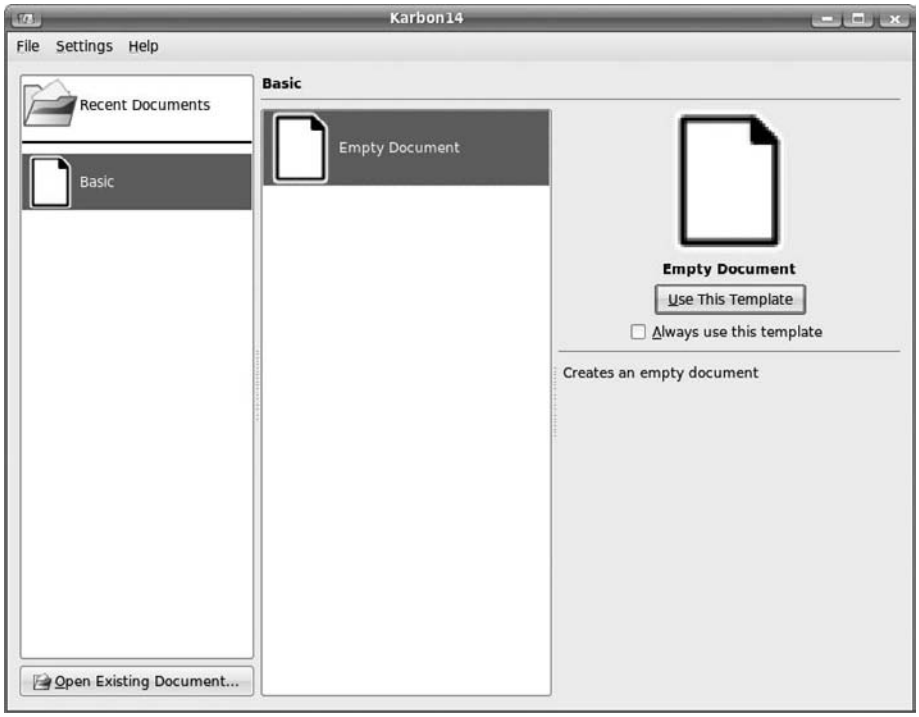
Starting Karbon14

To start Karbon14 on a Kubuntu system, select the K Menu ➪ Graphics ➪ Karbon14 - Scalable Graphics menu command. To start Karbon14 on an Ubuntu system, select the Applications ➪ Graphics ➪ Karb on14 menu item. You can also start Karbon14 by typing the karbon command in a GNOME Terminal, Konsole, or xterm window and pressing Return.

When Karbon14 starts, it displays its document wizard, as shown in Figure 18-24.

FIGURE 18-24

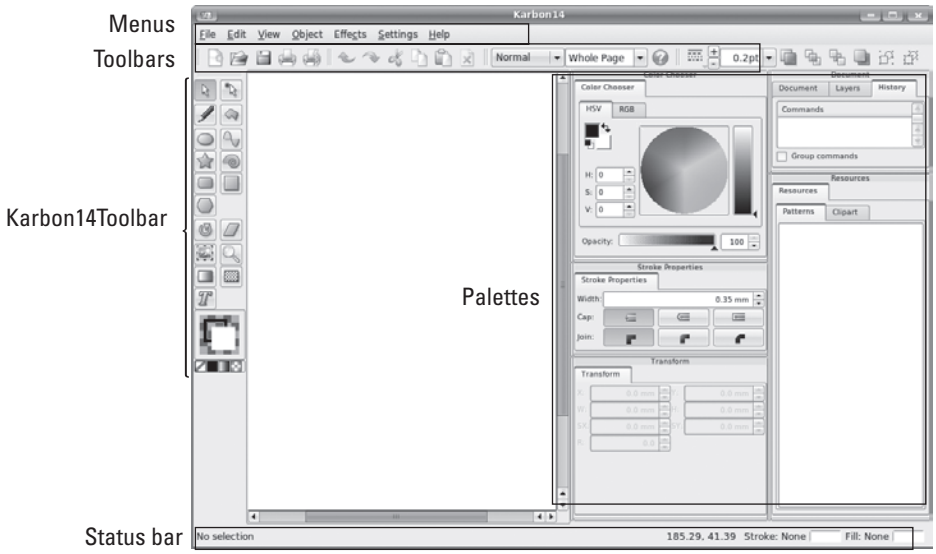
Karbon14's document wizard



Karbon14's document wizard enables you to select from various templates for any new image that you are creating (although it provides only a single template at the time of this writing), and also provides the Recent Documents button and Open Existing Document dialog to make it easy to re-open a Karbon14 document that you've worked on before. If you are creating a new image, click Use This Template to create a new, blank document. Karbon starts and displays a window like the one shown in Figure 18-25.

FIGURE 18-25

Starting Karbon14



NOTE **Depending on your screen size and resolution, Karbon14 doesn't always do a great job of displaying both your document and all of its default controls. You may find it necessary to use Karbon14 in full-screen mode, or to manually resize the Karbon14 window so that you can see a reasonable portion of the document that you are editing.**

A Quick Tour of Karbon14

Like all KOffice tools, Karbon14 provides extensive online help, so I won't bore you by walking through each menu and entry that it contains. However, as a quick overview of the Karbon14 interface, the callouts in Figure 18-25 highlight the following areas:

- **Menus:** The drop-down menus in Karbon14 are organized much like those in other drawing packages to provide a familiar user experience. Menu items with an arrow at their right edge lead to other, related submenus.
- **Toolbars:** By default, Karbon14 displays the File, Edit, View, and Object toolbars side-by-side below the menu bar. The vertical Karbon14 toolbar is shown to the left of the image editing area. You can modify the list of icons displayed on a toolbar by selecting the Settings ➪ Configure Toolbars command and selecting the toolbar that you want to edit at the top of the dialog that

displays. You can deactivate or activate individual toolbars by selecting Settings ➪ Toolbars and clicking the name of the toolbar that you want to deactivate or activate from the menu that displays.

- **Palettes:** Like other KOffice graphics tools such as Krita, Karbon14 displays specialized tool dialogs known as palettes to the right of the image editing area, each of which provides one or more tabs for controlling specific aspects of the image that you are creating or editing. Although they may be shown in a different order than that shown in Figure 18-25 when you open or create an image, these palettes are as follows:
 - **Color Chooser:** This palette displays HSV and RGB tabs to enable you to set the current fill and drawing colors and their opacity.
 - **Stroke Properties:** This palette displays a Stroke properties tab that makes it easy for you to quickly change the width of the current drawing stroke, and also enables you to specify the cap used at the end of a stroke (i.e., the shape used at the end of a stroke), as well as the corners used to join multiple strokes.
 - **Transform:** This palette displays a Transform tab that makes it easy for you to change the X and Y coordinates, height and width, rotation, and X and Y shear/skew for any selected portion of your graphic.
 - **Document:** This palette contains a Document tab that provides general information about the image that you are creating or editing; a Layers tab that enables you to show, hide, and lock any layer of your graphic by clicking on the icons in the entry for that layer; and a History tab that lists all of the changes that you've made to your graphic in the current editing session.
 - **Resources:** This palette contains a Patterns tab and a Clipart tab that make it easy for you to install (and subsequently identify) external clip art.

 You can show or hide the Palettes window or the tabs in different palettes by selecting or deselecting them using the commands on the View ➪ Palettes submenu.
- **Status Bar:** Shown at the bottom of the Karbon14 window, the status bar provides information about the current selection, image size, stroke, and fill settings.

Karbon14 provides extensive online help, available by selecting Karbon14 Handbook from the Help menu, or by pressing the F1 key on your keyboard at any time. Figure 18-26 shows the main page for the Karbon14 handbook.

Unfortunately, the Karbon14 documentation hasn't quite kept up with the development of the application, and thus doesn't discuss all of the latest features. If you become a Karbon14 fan, helping with the documentation would be an excellent way to contribute to the project. You can find contact information for the members of the Karbon14 team at `www.koffice.org/karbon/more_info.php`.

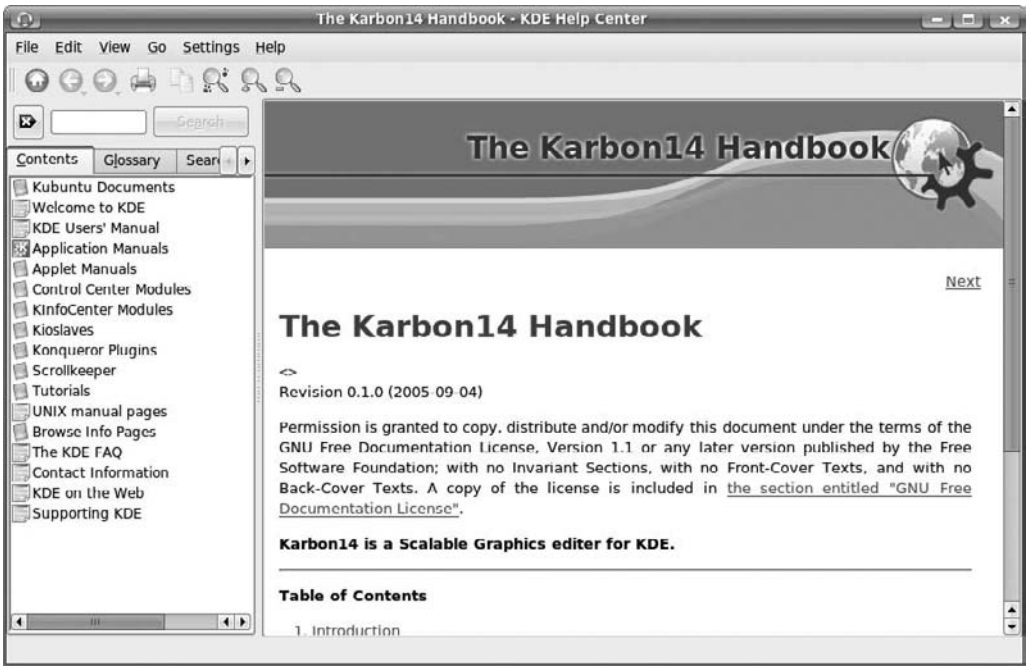
More Information About Karbon

The KOffice suite has been popular for quite a while on KDE-based systems. Karbon14 has been under development since 1998, but there are relatively few online resources available to provide information about using it. The primary Karbon14 site on the Web is `www.koffice.org/karbon`, which is a good starting point for information about Karbon14, its basic capabilities, and so on.

As Karbon14 matures and improves, it will certainly get more and more attention in the Linux and graphics press. A quick Clusty or Google search will help you find the latest and greatest information about Karbon14.

FIGURE 18-26

Karbon14's online help



Other Graphics Applications

Graphics programs seem to rank second to text editors and word processors in terms of the sheer number of available choices. Each has its own niche, fans, and detractors. Since I've already acknowledged that I'm not particularly good at graphics, I've covered only the graphics programs that I have actually used at one time or another and which I feel are the best applications for creating and editing various types of graphics on Linux systems. Some other good graphics applications are as follows (all of which are available in the Ubuntu repositories unless otherwise noted):

- **Blender:** An impressive 3D graphics modeling, rendering, and animation package that has been used to make various animated films, video and desktop computer games, and so on. Blender is also a great tool for architectural or similar wire-frame drawings. (`www.blender.org`)
- **CinePaint:** An image-editing application that began life as a clone of an earlier version of GIMP, but which has had a significant amount of development since then, including adding support for 8-bit, 16-bit, and 32-bit color channels and high quality image formats such as DPX, 16-bit TIFF, and OpenEXR. CinePaint has been used in a number of feature films (the "Harry Potter" films, *The Last Samurai*, the "Lord of the Rings" films, and so on), so it's impossible to dispute its quality or its business value. (`www.cinepaint.org`)

NOTE **CinePaint is not available in the standard Ubuntu repositories, but is included in the multimedia-oriented Ubuntu Studio Linux distribution (`http://ubuntustudio.org`).**

- **Dia:** A GNOME-based application for creating and editing diagrams and flowcharts on Linux systems, similar to and evocative of Microsoft's Visio application. Dia is not installed by default on Ubuntu systems. (`http://live.gnome.org/Dia`)
- **GNUPaint:** A simple, GNOME-based application for creating and editing graphics, GNUPaint has all of the attributes of xpaint (its ancestor), Microsoft Windows Paint, or MacPaint. GNUPaint is easy to use, but that's primarily because it doesn't do much when compared with the other graphics applications discussed in this chapter. GNUPaint is not available in the standard Ubuntu repositories. I'm primarily mentioning this here because you might have heard of it, not as a suggestion. (`www.gnu.org/software/gpaint`)
- **KChart:** A KDE-based application for creating and editing charts. KChart is a component of the KOffice suite of applications, and makes it easy to create charts that can be embedded in KWord documents, KPresenter presentations, and so on. KChart is not installed by default on Kubuntu systems. (`www.koffice.org/kchart`)
- **Kivio:** A KDE-based application for creating and editing diagrams and flowcharts, once again similar to and evocative of Microsoft's Visio application. Kivio is a part of the KOffice suite of applications, although it is not installed by default on Kubuntu systems. (`www.koffice.org/kivio`)
- **Xara Xtreme:** A vector graphics application that was primarily developed for Microsoft Windows systems (and is still actively supported and under development there), Xara Xtreme was released as open source in 2006 and ported to Linux. This is an excellent, fast, and very usable graphics application that I've used numerous times. Its creators deserve significant kudos for releasing a commercial app as open source, but their efforts didn't capture the fancy of a substantial developer community (for both good reasons and bad). Xara Xtreme is not available in the standard Ubuntu repositories, and its Web site (`www.xaraxtreme.org/about/`) hasn't been updated for a while. I suspect that using Xara Xtreme on Linux systems is an evolutionary dead end unless you want to maintain it yourself. And if you do, please let me know, because I really liked this application. Alas, poor Xara Xtreme, I knew thee well…

This list, and the applications that were discussed at length in this chapter, should help you create and edit almost any type of graphical images or objects on an Ubuntu or Kubuntu system. It's quite possible that I've missed some other excellent graphics applications for Linux — if you fall in love with some application that isn't mentioned here, please let me know. Inquiring minds want to know!

Summary

This chapter explored tools for creating and manipulating various types of graphics files on Linux systems. Linux has a great selection of powerful graphics tools that you can use to do real work that is constrained only by your artistic skills. (I am still waiting for a "mental image to graphics file" application, but that's my problem.) GIMP, Krita, OpenOffice.org Draw, Inkscape, and Karbon are all shining examples of just how excellent and powerful open source software can be.

Chapter 19 explores how to work with audio and video on Ubuntu and Kubuntu systems, with topics ranging from how to configure sound on your system, to how to work with CDs, DVDs, and the alphabet soup of digital audio and video file formats that one finds on the Net today.

Chapter 19

Working with Multimedia

MP3, WMA, OGG, FLAC, SHN, APE: this may look like a blast from a shotgun loaded with three-letter acronyms (TLAs), but it's really a list of the most popular formats for digital audio nowadays. The same is true of digital video, where acronyms such as MPG, MPEG, and VOB are king. There are also some really ugly acronyms such as MPAA and RIAA, which stand for groups of self-interested thugs trying to deny the rights of the consumer, but we won't go into that here.

The great thing about digital audio and video is that you no longer need specialized devices to play your music or watch a movie. You can listen to a music CD or play a video CD or DVD on your computer. Once you buy a CD or DVD, it is completely legal for you to transfer its contents to your computer and watch it there. This eliminates the need to carry around the original media, and you can safely leave them behind at home as backups in case your computer crashes.

This chapter helps you make some sense out of the maze of acronyms, general terminology, and specialized jargon that aficionados use to describe digital audio and video, and explains how to use the software available for Linux to listen to music, watch movies, and even create your own CDs and DVDs. I love being able to check into a hotel room and relax while listening to my favorite tunes because I've brought them along on my favorite modern multimedia device—a computer running Ubuntu or Kubuntu Linux.

IN THIS CHAPTER

- Understanding multimedia terminology
- Configuring sound on Ubuntu
- Configuring sound on Kubuntu
- Exploring the Linux multimedia stack
- Playing, ripping, and burning CDs and DVDs
- Working with audio files and Internet radio

Overview of Digital Audio and Video Terminology

As mentioned in the introduction to this chapter, digital audio and video is an area in which the computer industry's tendency towards using jargon and three-letter acronyms makes many conversations and reference information almost incomprehensible to the novice. Though I've been a record collector, CD

collector, and musician for years, I still encounter terms that I either have to look up or decrypt, and there is no single source of this sort of terminology outside of all-encompassing sites like Wikipedia.

This section, therefore, provides a basic glossary for the most common audio, video, and digital multimedia terminology and acronyms that I am familiar with and see often. It isn't necessarily easy reading, but it should be a great reference for the rest of the chapter, where these terms will be used.

- **AAC (Advanced Audio Coding):** A lossy codec originally developed by Dolby Laboratories that is similar to MP3 files in file size but provides slightly higher quality. AAC is now often, and incorrectly, referred to as the Apple Audio Codec because of Apple's use of a protected version of the AAC format to impose DRM on its iTunes customers.
- **ABR (average bit rate):** A type of encoding in which the user specifies the target bit rate for sampling and playback, and the encoder calculates the appropriate VBR (Variable Bit Rate) values necessary to achieve the average.
- **AIFF (Audio Interchange File Format):** A lossless audio codec originally developed by Apple Computer, and therefore commonly used on Macintosh systems.
- **ALAC (Apple Lossless Audio Codec):** A proprietary lossless audio code developed by Apple Computer. This format is frequently confused with the M4A format, which is essentially a container file that can deliver files in ALAC and other formats.
- **APE:** The file extension for audio files encoded using the proprietary lossless Monkey's Audio codec. This codec provides slightly better compression (i.e., slightly smaller file sizes) than the FLAC or SHN codecs, but at the cost of slower encoding/decoding. APE files are primarily used for lossless compression and archival on Microsoft Windows platforms because no official Linux port of the Monkey's Audio codec exists. For more information about the Monkey's Audio codec, see `www.monkeysaudio.com`.
- **bit rate:** The number of bits per time interval that make up an audio signal. The most common time interval used in bit rate calculations is one second (BPS). High-speed data communications is often measured in Kbps (kilobits—thousands of bits per second). This is overkill for digital audio.
- **BitTorrent:** A protocol for downloading and uploading digital data, including audio and video files. BitTorrent isn't really a multimedia term, except that it provides a common mechanism for sharing multimedia files, and I thought you might look for it here. See the section of Bonus Chapter 2 entitled "Using BitTorrent" for information on the applications that you can use to download and upload files using the BitTorrent protocol on Ubuntu and Kubuntu systems.
- **burn:** The term for writing data to recordable or rewritable CD or DVD media. The term *burn* comes from the fact that data is written to CDs and DVDs using a laser that inscribes miniscule pits on their surface to represent digital information.
- **CBR (constant bit rate):** A codec that uses a fixed number of bits per second as opposed to ABR or VBR codecs.
- **CD (compact disc):** A storage device capable of holding approximately 650MB of prerecorded/prewritten digital information.
- **CD-R (compact disc/recordable):** A storage device capable of holding approximately 650MB of digital information, but where each portion of the disc can be written to once by home computers with compatible drives. Writing a CD-R disc can be done in multiple sessions, where each session writes a different portion of the disc. Mike Richter's Primer on CD-R at `www.mrichter.com/cdr/primer/primer.htm` is a tremendous resource for everything you ever wanted to know about CD-R (and more). Similarly, Andy McFadden's CD-Recordable FAQ at `www.cdrfaq.org` is another tremendous resource with tons of detailed information about CD software. Web

documents like these are true testimonials to the power of the Internet and the expertise and dedication of its users. My hat goes off to these guys!

- **CD-RW (compact disc/rewritable):** A storage device capable of holding approximately 650MB of digital information, which can be written to multiple times by home computers with compatible drives. Only CD-RW drives can read CD-RW discs.
- **CDDA (Compact Disc Digital Audio):** The format used on most non-copy-protected audio CDs; sometimes also referred to as CDA because of Microsoft Windows' inability to use four-character file extensions.
- **CDDB (CD Database):** A commercial, online database that provides identification information and track listings for commercial CDs.
- **codec (co**mpressor/dec**ompressor):** An algorithm by which digital signals are compressed so as to lose a minimum amount of quality, and decompressed for playback.
- **DAT (digital audio tape):** A digital recording tape medium used in recording studios and by audio professionals.
- **DAE (Digital Audio Extraction):** The process of cloning prerecorded digital media by directly copying it to files rather than by playing it, converting the resulting analog audio output, and saving that information to a digital storage device.
- **disc-at-once:** A CD/DVD burning mechanism that writes the entire surface of the CD or DVD disc rather than simply the amount of storage required to hold the information that you are recording.
- **DMCA (Digital Millennium Copyright Act):** Legislation passed in the United States in which the supposed goal was to preserve copyrights for digital media but that has been used to squash reverse engineering, eliminate free-speech discussions of encryption technology, and rob users of the free reuse of legitimately purchased audio and video files; a good idea gone bad.
- **DRM (Digital Rights Management):** A technique for denying legitimate users of audio and video files the ability to freely copy and use them.
- **DSP (digital signal processing):** Using various algorithms to convert or modify digital audio and video signals. DSP can be done in hardware, software, or a combination of both.
- **DVD (digital video disc):** A storage device capable of holding approximately 4.7GB of digital information.
- **EQ (equalize):** A recording effect that increases or decreases the volume of different frequency ranges in an audio signal. Equalizing frequency ranges smoothes the audio signal by preventing extreme differences in volume across different frequency rangers.
- **FFP (FLAC footprint):** Checksum files that are typically distributed with FLAC-encoded files to enable you to verify their integrity. See the entry for MD5 in this glossary for more information about checksums.
- **FLAC (Free Lossless Audio Codec):** A lossless audio codec that typically achieves compression rates of 30% to 50% without any reduction in quality from the original recordings.
- **Freedb (Free DB):** An open source alternative to CDDB that enables users to contribute information about CDs.
- **ID3 tags:** 128-byte header information in MP3 files that is used to provide information about the audio file. ID3v1 tags provide album, artist, general comments, genre, and title information. ID3v2 tags provide this information and more in a more free-form fashion, permitting larger, user-defined tags that support all ID3v1 tags but can also include information such as lyrics and images.

- **ISO (International Standards Organization):** In the context of CDs and DVDs, ISO refers to the three-letter file extension added to CD and DVD images to show that the contents of the file conform to the ISO9660 specification, which defines CD and DVD formats. People often refer to conformant image files that contain complete CD or DVD images as ISOs.
- **lossless formats:** File formats that are not lossy, that is, file formats that provide all of the information that was present in the audio, graphics, or video information as it was originally recorded, drawn, or captured.
- **lossy formats:** File formats that sacrifice quality to reduce file size. The term *lossy* refers to the data compression algorithm used to reduce file size, whereby some of the information that was present in the original file is actually eliminated to produce the smaller, compressed file, and therefore the original file can never be completely reconstructed from the compressed version. In the case of audio files, lossy images typically reduce both the bit rate and the frequency ranges at which audio is recorded, the theory being that the majority of this information is inaudible to most people.
- **M4A:** An audio file container format that can contain AAC or (more commonly) ALAC files. Protected versions of M4A files that have been subjected to Digital Rights Management are delivered as **M4P** files.
- **MD5:** A popular mechanism, known as a *checksum*, used to verify the integrity of digital data files, especially lossless audio files. Files with this extension are plain-text files that typically contain unique 32-character ID numbers for one or more companion data files. An MD5 checksum value is unique to the data file that accompanies it, and therefore is often used to fingerprint (i.e., uniquely identify) that file and thus verify its integrity. If you generate an MD5 checksum (using a Linux application such as md5sum) and that checksum value matches the checksum value delivered with the file, the file has not been modified.
- **MP3 (MPEG-1 Audio Layer 3):** The most popular audio format used today, MP3 is a proprietary lossy codec that produces relatively high-quality sound from relatively small audio files.
- **MP3 CD:** A compact disc that contains MP3 files rather than traditional CDDA audio. Because MP3 files are smaller than CDDA files, MP3 CDs can hold much more music than other CD formats, and are quite popular nowadays even though they usually cannot be played in standard CD players.
- **MPAA (Motion Picture Association of America):** The motion picture industry's equivalent of the RIAA, a group of industry thugs bent on maximizing industry income at the expense of the end user and often the artist.
- **MPEG (Motion Picture Experts Group):** An Industry consortium responsible for defining standards for audio and video formats.
- **MPEG-4:** The latest generation of MPEG standard at the time that this book was written, MPEG-4 is a standard file format for delivering audio, video, and other multimedia objects.
- **noise reduction:** Audio processing techniques used to clean up audio files by removing, filtering, or equalizing specific frequency bands.
- **normalize:** An audio processing technique that identifies the loudest range of an audio file, and then uses that as a baseline for the volume used when playing back the rest of the file. Normalizing audio data does not change relative volumes within a song, just the baseline volume.
- **OGG (Ogg Vorbis):** A popular, open source lossy codec and associated container format that provides a powerful, royalty-free alternative to lower quality codecs such as MP3. Ogg Vorbis was developed by the Xiph.org foundation, a foundation dedicated to developing and promotion free, open source multimedia formats. (Ogg is the file format, while Vorbis is the actual codec, though

files encoded using this codec typically have the .ogg file extension.) For more information about Ogg Vorbis, see `www.xiph.org` or `www.vorbis.com`. If your software or audio player does not support Ogg Vorbis, contact the manufacturer.

- **P2P (Peer-to-Peer):** A popular data exchange and delivery mechanism where each client (peer) identifies other peers and exchanges data with them without the use of a central server. Popular P2P applications are BitTorrent, GNUtella, LimeWire, and so on.
- **PCM (Pulse Code Modulation):** An audio and video data format that provides a true representation of the recorded signal by sampling it at regular intervals. PCM is the standard form for digital audio in computers, compact discs, and in digital video.
- **RealAudio:** A popular streaming audio format developed by Real.com and often used on the Web. For more information, see `www.real.com`.
- **RIAA (Recording Industry Association of America):** A group of industry thugs bent on maximizing industry income at the expense of the end user and often the artist.
- **rip:** The generic term used to describe the process of converting CD and DVD discs to files that can be used on a computer system.
- **SHN (Shorten codec):** A popular lossless audio codec frequently used for distributing authorized live recordings on the Internet. Shorten delivers slightly larger file sizes than FLAC, its primary "competitor." Interestingly, the Shorten codec also offers a lossy mode, although I don't see why anyone would use that. For more information about the Shorten codec and Linux software for working with SHN files, a great resource is the SHN Resources page at `http://research.umbc.edu/~hamilton/shnlinks.html`. The Shorten codec is a commercial product of SoftSound Ltd., but has a user-friendly personal use license. For more information, see `www.softsound.com/Shorten.html`.
- **skin:** An alternate user interface for graphical applications. New skins can be applied to an existing application without recompilation or any actual changes to the application other than its configuration information. Many audio players (such as XMMS and the venerable WinAmp for Microsoft Windows systems, which pioneered skinnable applications) are skinnable, which enables fans to easily customize their appearance.
- **streaming:** A method of delivering digital content over a network such that an application (such as an audio player) receives information as it is needed without the need to store that data on the machine. The most common examples of streaming media players are the RealPlayer and RealOne applications from RealAudio. For more information, see `www.real.com`.
- **SVCD (Super Video CD):** The successor to Video CDs (VCDs), an SVCD contains MPEG-2 video and MPEG-1 or MPEG-2 audio recorded on a normal VCD. Most DVD players can also play SVCD discs because DVD players support the MPEG-2 codec required to decode the MPEG-2 video data.
- **track-at-once:** A CD/DVD-burning mechanism that writes on track of a CD or DVD disk but leaves the remainder of the CD or DVD available for further recording at a later stage.
- **transcode:** The process of converting from one encoding format to another, typically used when referring to converting a media file or object from one format to another, as is done when creating DVDs from MPEG files in other formats (or vice versa).
- **TTA (True Audio codec):** A free lossless audio codec offering an average of 30% to 40% file size reduction without any loss of quality. I haven't encountered this codec much on the Net, but some people swear by it. For more information about the True Audio codec, see `www.true-audio.com`.
- **Torrent Files:** The files that contain information about how to contact a BitTorrent server in order to upload or download digital files using the popular BitTorrent protocol. "Torrent Files"

isn't really a multimedia term, except that BitTorrent provides a common mechanism for sharing multimedia files and I thought you might look for it here. See the section of Bonus Chapter 2 entitled "Using BitTorrent" for information on the applications that you can use to download and upload files using the BitTorrent protocol on Ubuntu and Kubuntu systems.

- **VBR (variable bit rate):** A type of encoding that attempts to reduce file size without sacrificing significant quality by varying the bit rate at which audio information is recorded. More complex audio is recorded at a higher bit rate, while simpler audio passages are recorded at a lower bit rate.
- **VCD (Video CD):** A storage device capable of holding approximately 650MB of prerecorded/prewritten digital video and audio information that are encoded in MPEG-1 format. Most DVD and CD players can play VCDs because they support the MPEG-1 codec used to decode the video and audio data.
- **WAV:** A digital audio file format originally developed by Microsoft and IBM that has become the standard format for distributing lossless audio (waveform) files. WAV files consist of chunks of data that can actually be encoded using any codec but most commonly contain raw pulse-code modulation (PCM) audio data, which is lossless and therefore huge. WAV files are often used by audio professionals to ensure maximum quality and complete fidelity to the original recordings.
- **Wavpack:** A free lossless audio codec used primarily on Windows systems but also available for many different Linux distributions, architectures, and applications such as XMMS. I haven't encountered Wavpack much outside of the Windows community, but your mileage may vary. Wavpack is extremely quick at compressing and decompressing files with no loss of quality. For more information about Wavpack, see `www.wavpack.com`.
- **WMA (Windows Media Audio):** A popular codec for Microsoft Windows users that is the default audio codec used by the Windows Media Player. Lossy, lossless, and DRM versions of the WMA codec are available. The most important thing that Linux users need to know about WMA is how to convert audio files encoded with this codec into more reasonable formats.
- **WMV (Windows Media Video):** A popular lossy codec for Microsoft Windows users that is the default video codec used by the Windows Media Player. The most important thing that Linux users need to know about WMV is how to convert video files encoded with this codec into more reasonable formats.

Unless you have plenty of free brain cells and nothing else to do with them, there's no need to memorize these terms—that's what books like this one are for. However, now that you have expanded your digital multimedia terminology, let's move on to something more important—configuring the sound devices on your system so that you can actually play audio and video on your Ubuntu or Kubuntu system. Since Ubuntu and Kubuntu use quite different sound systems, they are discussed separately in the next two sections.

Configuring Sound Devices, Levels, and System Sounds on Ubuntu Systems

Most of the setup for your Ubuntu system's audio components is done automatically when you install your Ubuntu Linux system. Ubuntu's installer has excellent hardware detection capabilities and automatically installs and configures the ALSA (Advanced Linux Sound Architecture) and PulseAudio utilities, drivers, and libraries that Ubuntu uses to provide excellent, high-quality audio playback. The following sections enable you to specify and test the sound devices used by your system, set sound levels on your

system, and customize the sounds your system plays in response to various system events, and tell you how to set preferences for CD and DVD handling.

Specifying and Testing Sound Devices

To verify that sound is working on your Ubuntu Linux system, first check the obvious things, such as whether your speakers are plugged into your sound card's audio out port (usually a light green socket—the central one—of the three sound ports on the back of your system or on your sound card), the speakers are turned on, and so on. Next, select the System ➪ Preferences ➪ Sound command. The Sound preferences utility appears, as shown in Figure 19-1.

FIGURE 19-1

Default system sound settings in the Sound Preferences utility



As you can see from this dialog, most of the default devices associated with different types of audio on your system are initially set to Autodetect, which means that they use whatever devices the system believes it should use for those events. To test any of these, click the Test button to the right of the device associated with a certain type of sound playback. A dialog entitled Testing Pipeline appears, and you should hear a sound until you click OK in that dialog. (The term *pipeline* refers to the series of applications that are used to test sound playback of a given type.)

If you do not hear a sound when testing a certain type of sound playback, click the drop-down list to the left of the Test button. This displays the list of available audio devices and sound systems that are available on your system, as shown in Figure 19-2.

FIGURE 19-2

Drop-down showing available audio devices and sound systems



The audio devices that you see are specific to your hardware, and the available sound systems will be the following:

- ALSA - Advanced Linux Sound Architecture
- ESD - Enlightened Sound Daemon (Ubuntu 8.04 only)
- OSS - Open Sound System
- PulseAudio Sound Server

Auto-detection can be a problem when your system contains multiple audio devices, and especially if you are using an USB audio device, since the system's sound detection setup routines usually try to use dedicated sound hardware, such as an integrated sound device or a sound card, before trying generic sources to which audio devices may be attached. (I'll discuss the available sound systems later in this chapter in the section entitled "Exploring the Linux MultiMedia Stack".)

If you did not hear a sound for a specific type of sound playback, select one of the available sound devices and then click Test. If you hear a sound, great—if not, repeat this process until you find an audio device that works.

TIP **Most motherboards nowadays feature integrated audio hardware. Unless you've added a sound card to your system (in which case you probably know what it is), you can identify the primary audio device on your system by using the `lspci` command, as in the following example:**

```
$ lspci | grep -i audio
00:1b.0 Audio device: Intel Corporation 82801H HD Audio...
```

Of course, you'll see a different audio devices on your system, but the entry that is displayed should at least give you an idea of which of the available audio devices to try selecting first. To see device-level information about your sound hardware, see the section later in this chapter entitled "Getting Detailed Information about Your Sound Hardware."

The Default Mixer Tracks section of this tab identifies the sound device, associated mixer application, and specific track of that sound, which you will be able to control using the arrow keys or any audio control buttons on your keyboard. On modern Ubuntu systems, this will almost always be one of the choices that use the ALSA Mixer application, because local sound output is initially processed using ALSA, or PulseAudio if you are using a sound server. The list of available mixer tracks depends on the sound device or server, but you will usually want to select either Master (the primary software volume control), Volume, (the primary software volume control for some sound devices), or PCM (the actual audio device input/output level). Master (or Volume) is the most common track to select, but I have found it useful to select PCM here and then use the panel's Volume Control applet (which I'll discuss later in this section) to control the overall audio volume. This provides keyboard access to the two primary volume levels that most people worry about.

TIP **If you want to be able to control multiple sound sources with the arrow keys or your keyboard's volume key, you can hold down the Ctrl key and select multiple sources. This is especially useful if your audio device offers multiple PCM settings.**

Setting System Sound Levels

The section of Chapter 5 entitled "Customizing Panel Contents" discussed how to add various applets to your GNOME desktop's panels. One of the more useful applets that is preinstalled for you in the upper panel on your Ubuntu Desktop is the Volume Control applet, whose icon looks like a small speaker and is usually displayed at right in the top panel, just to the left of the date.

To customize your system's volume settings for the current session, right-click on this icon. The slider shown in Figure 19-3 is displayed.

To change your current volume settings, drag the slider up or down to the value that you prefer.

To get access to more detailed volume settings, left-click the volume control icon and select the Open volume control command to display the dialog shown in Figure 19-4.

NOTE **The volume settings displayed in the dialog shown in Figure 19-4 are based on the capabilities of your sound card. Based on the sound card that your system is using, you may see additional settings.**

If your system has multiple sound cards (or a built-in sound card and a standalone sound card that you installed yourself), make sure that the sound card you are changing settings for is the card to which your speakers are attached and that is listed in the header of the Volume Settings dialog shown in Figure 19-4. To change the sound device that you are configuring, select the Devices drop-down list at the top of this dialog and select the correct sound interface.

FIGURE 19-3

Quickly changing current volume settings



FIGURE 19-4

Detailed playback volume control settings



TIP **On Ubuntu 8.04 systems, you will have to select the File ➪ Change Device menu command and select the appropriate device from the drop-down list there.**

The default playback volume settings that you can adjust for most sound cards are **Master**, the volume setting for all audio playback from any sound-enabled application, and **PCM**, the volume setting for the audio card itself. You may also see other settings, such as **Front** (shown in Figure 19-4), which is available if your system has additional audio outputs on the front of the system.

Two icons are displayed below each slider:

- The chain icon determines whether the two sliders for each output (the left and right channels) must move together or can be set separately. A linked chain means that the channel sliders will move together. A broken chain indicates that they can be set separately. Click the icon to toggle whether the sliders are linked together or not.
- The speaker icon determines whether the output is muted or not. If a red X is displayed over the speaker icon, that output is currently muted. Click the icon to toggle whether an output is muted or not.
- You can show additional settings and tabs in this dialog by clicking Preferences (or selecting Edit ➪ Preferences in Ubuntu 8.04) and selecting the checkboxes associated with other sound settings in the dialog that displays. For example, to change the volume settings associated with capturing/recording sounds, select the checkbox beside the Capture entry. This will display a Recording tab in the main Volume Control dialog, which contains the control(s) for modifying input sound levels. In addition to the icons displayed below the sliders on the Playback tab, the Recording tab also displays a Microphone icon, which indicates that recording is available. If a red X is displayed over this icon, that input source is currently disabled. Click the icon to toggle whether an input/capture source is active.

On Ubuntu 8.10 systems, you can also change the ports on your sound card that are used as output ports by selecting the Switches tab and selecting the checkboxes beside the ports that you want to use as output devices.

Once you have modified any volume settings that you want to change in the tabs of the Volume Control applet, select File ➪ Quit to close the Volume control applet.

Testing and Customizing System Sounds

By default, your system is configured to play a "welcome" sound when you log in. To configure (or disable) the sounds that your system plays in response to various events, select the Sound Preferences utility's Sounds tab. The dialog shown in Figure 19-5 displays.

NOTE **This dialog looks somewhat different in versions of Ubuntu prior to 8.10 but provides the same capabilities. Earlier versions provide a slightly simpler interface in which the sounds associated with various system events are listed in a drop-down menu. This menu contains a No Sound entry (which disables playing a sound for that event), various predefined sounds, and a Select sound file entry that enables you to associate a custom sound file with that even. To change the sound associated with a system event, select No sound, another predefined sound, or your own sound file via the Select sound file entry. You can click the Play button to the right of this menu to test your new sound/event association.**

FIGURE 19-5

Default system sound settings in the Sound Preferences utility



The Sound Preferences utility enables you to verify the system activities for which your system will play an associated sound. The settings for each system event with which a sound can be associated are the following:

- **Default:** Play the default sound associated with an event when that event occurs.
- **Disabled:** Do not play a sound when the system event occurs.
- **Custom:** Play a user-specified sound file when a system event occurs.

To change, disable, or enable the sound associated with any system event, left-click on the current sound setting for that event. This displays a pop-up menu from which you can select any of the three settings in the previous list.

TIP **To be able to associate sounds with button clicks and window events, you will have to check the Play sound effects when buttons are clicked checkbox at the top of the screen.**

To play the sound that is currently associated with any system event, click on the right-arrow button to the right of any system action that currently has a sound associated with it (i.e., is not marked as Disabled). If you hear a sound, congratulations! Sound is working correctly on your Ubuntu system.

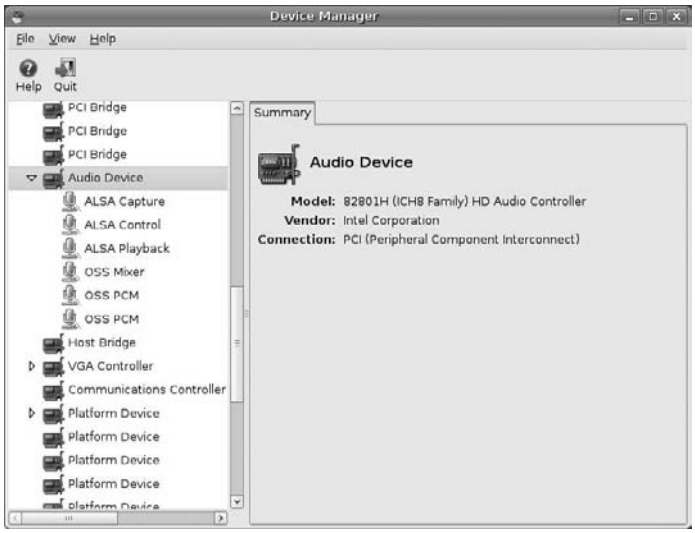
Getting Detailed Information about Your Sound Hardware

If you've gotten to this point and sound still isn't working on your Ubuntu system, it may be time to explore your system's sound hardware to see if it is set up correctly and to give you the information that you need to ask for help. Information about your system's hardware, including your sound hardware, is available in Ubuntu's Device Manager application, which is not installed by default but can be quite handy when looking for detailed information about your system's hardware. To install this application, use any of the software installation tools discussed in Chapter 23 to install the `gnome-device-manager` package.

The GNOME Device Manager, which you can display by selecting the System ➪ Administration ➪ Device Manager menu command or executing the `gnome-device-manager` command from any command-line, is shown in Figure 19-6.

FIGURE 19-6

Sound hardware entries in the Device Manager



The entries in the Device Manager for your sound hardware will tell you whether Ubuntu has been able to identify the device, what it has identified it as, and whether valid entries for that device exist in the /sys filesystem, which is a virtual filesystem that provides direct access to various aspects of your system hardware.

TIP **If your system sounds play correctly, but you can't play CDs, make sure that your CD player is correctly connected to your sound hardware. This connection typically consists of two or three pin cables that run from the back of your CD drive to your sound card or the sound hardware that it built into your motherboard. Some sound cards and motherboards support several input sound devices, and it's quite possible that your CD drive is connected to the wrong port or that the connector has become dislodged.**

If the entries in the Device Manager appear to match your hardware, it's time to search the Web, starting with various Ubuntu resources listed in the section of Chapter 1 entitled "Support for Ubuntu Linux." Paid support is always an option, but your best (and certainly cheapest) bet is to begin by searching the various forums for information about your sound hardware or the specific problem that you are seeing. If you don't find any relevant or helpful information, try posting to the Ubuntu Users mailing list. As emphasized throughout this book, one of the best aspects of Ubuntu Linux is its huge, helpful user community. You're also welcome to ask me, but my response time is certainly going to be slower than that of the thousands of continuous participants in the Ubuntu forums and mailing lists.

Setting Blank CD and DVD Preferences

Now that few people bother with floppy disks (or, perhaps, know what they are), CDs and DVDs are the most common examples of removable media that you'll use with your computer system. Removable media is the generic term for any media that contains digital data and that is not a mandatory part of your computer system. Removable disks, flash drives, and other devices that show up as disk drives (like an iPod) are not exactly the same things as far as Ubuntu is concerned, and are discussed in Chapter 21, "Consumer Electronics and Ubuntu," as well as in Chapter 24, "Adding Hardware and Attaching Peripherals." This section focuses on how you can configure what your system does when you insert a prerecorded or blank CD or DVD.

When you insert a blank CD or DVD in your Ubuntu system, the system does two things:

- It displays an appropriate icon for the type of blank media that you have inserted.
- It displays a dialog asking what you would like to do with that media.

Figure 19-7 shows the icon and dialog displayed when you insert a blank CD-R or CD-RW disc in a compatible drive on your system. A very similar icon and dialog are displayed when you insert a blank DVD-R disc into a compatible drive on your system.

FIGURE 19-7

Dialog and icon for blank CD-R/RW insertion



In either case, clicking Cancel closes the dialog without starting an application. Clicking OK opens Ubuntu's CD/DVD Creator application, which is the Nautilus file manager in a special mode. I prefer to click the drop-down menu and select Brasero, a specialized application for creating CDs and DVDs, which is discussed later in this chapter. If you always want to open Brasero when you insert a blank CD or DVD, you can select the Always perform this action checkbox before clicking OK.

How your system reacts when a blank CD or DVD is inserted, and the applications that are available in the dialog shown in Figure 19-7, are defined on the Nautilus Preferences dialog shown in Figure 19-8.

FIGURE 19-8

Setting preferences for blank CD/DVD or other media



To view this dialog, start Nautilus by clicking on an existing filesystem icon or by selecting an entry from the GNOME Desktop's Places menu. Select the Edit ➪ Preferences menu item, and click the Media tab in the dialog that appears.

NOTE **Versions of Ubuntu prior to 8.04 integrated these settings into a Removable Drives and Settings application that was available on the System ➪ Preferences menu. Because this information is actually part of your GNOME settings, newer Ubuntu distributions leverage the existing settings in GNOME's Nautilus file manager rather than providing and supporting a separate application.**

To prevent your system from displaying the dialogs shown in Figure 19.7, select Blank CD disc or Blank DVD disc from the drop-down Type menu in the Other Media section, and either select the Do nothing entry or select the application that you always want to open when blank media is inserted. You can specify a custom application by selecting the Open with other application menu entry and browsing for the location of the application that you want to use. (This is a great way to integrate my favorite CD/DVD burning application, k3b, into your Ubuntu system.)

How your Ubuntu system reacts when you insert a prerecorded CD or DVD depends on the preferences settings for CD Audio and DVD Video on the same tab. To disable automatically starting a specified application, click the drop-down menu beside the appropriate type of media and select Do nothing. To change the application that starts automatically when prerecorded media is inserted, click the drop-down menu beside the appropriate type of media and either select one of the applications that is listed on the menu or select Open with other application and browse for the location of the application that you want to use.

For information about playing or burning CDs and DVDs, see the sections later in this chapter entitled "Working with CDs" and "Working with DVDs."

Configuring Sound Devices, Levels, and System Sounds on Kubuntu Systems

Most of the setup for your Kubuntu system's audio components is done automatically when you install your Kubuntu Linux system. Kubuntu's installer has excellent hardware detection capabilities and automatically installs and configures the ALSA (Advanced Linux Sound Architecture) utilities, drivers, and libraries that Kubuntu uses to provide excellent, high-quality audio playback. The following sections enable you to specify and test the sound devices used by your system, set sound levels on your system, customize the sounds your system plays in response to various system events, and tell you how to set preferences for CD and DVD handling.

Specifying System Sound Devices

To verify that sound is working on your Kubuntu Linux system, first check the obvious things, such as whether your speakers are plugged into your sound card's audio out port (usually a light green socket—the central one—of the three sound ports on the back of your system or on your sound card), the speakers are turned on, and so on. Sound preferences on both KDE 3 and KDE 4 systems are set within the K Menu ➪ Systems Settings application. To begin configuring the sound devices on your system, select the Sound System configuration applet on KDE 3–based Kubuntu systems or the Sound applet on KDE 4–based systems. The next two sections explain how to configure sound on each of these types of systems because (of course), dialogs and options are slightly different—not only because of creeping featurism between KDE 3 and KDE 4, but also because the sound systems used on the two versions of KDE differ. KDE 3 uses the analog Real-time synthesizer (aRts) sound system (RIP), while KDE 4 uses its new Phonon sound system.

Configuring Sound on KDE 3 Kubuntu Systems

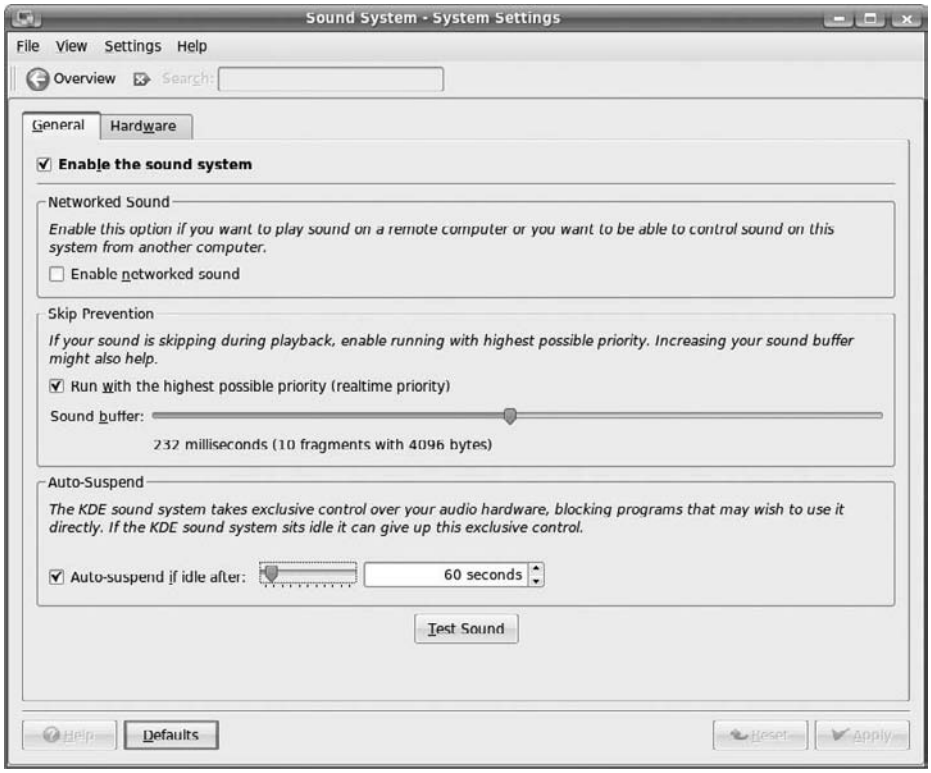
Figure 19-9 shows the default Sound System configuration screen on a KDE 3–based Kubuntu system.

The most important checkbox on this panel is the Enable the sound system checkbox, which must be selected for audio to work on your Kubuntu system. The Enable network audio checkbox is only important if you want to be able to play audio from your system on another system or if you want to be able to

control the audio on your system from another system. I rarely do this—if I want to play audio files that are stored on one system on another, I share the folder that contains them using Samba or NFS.

FIGURE 19-9

Default system sound settings in the KDE 3 Sound applet



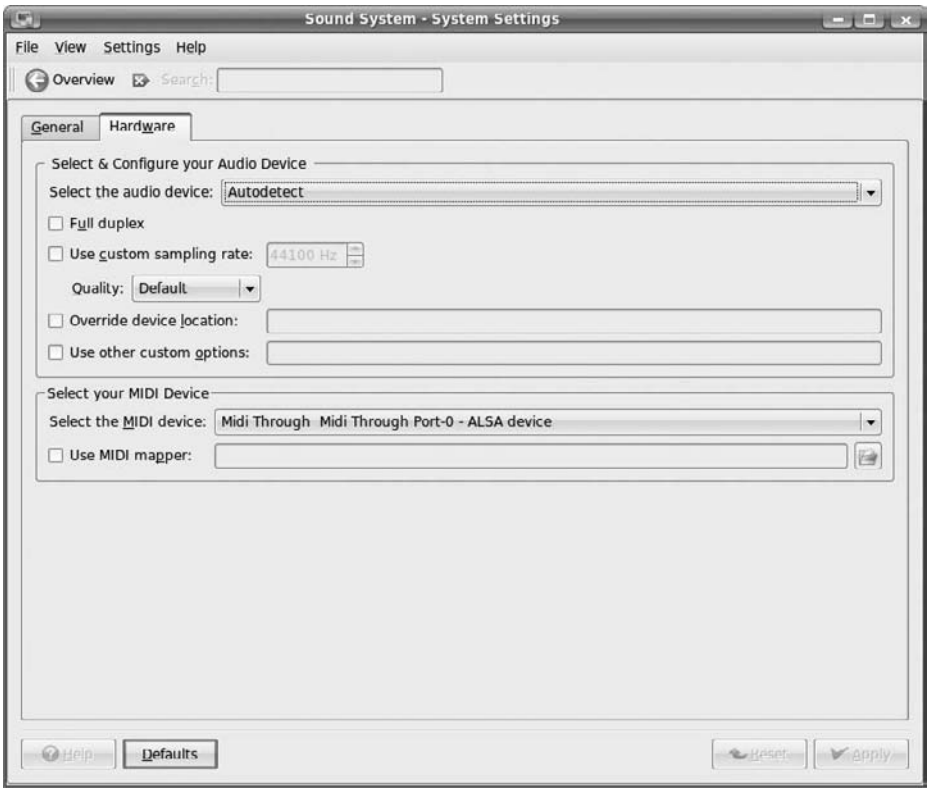
The Skip Prevention (which means "preventing audio from skipping," not skipping prevention) section can be very handy if your system is heavily loaded and sound playback is being interrupted by other applications or system events. The Run with the highest priority (real-time priority) checkbox is normally selected by default, but you can also increase the size of the buffer used to hold sound data before it is actually played if the audio on your system is still choppy under heavy loads.

Finally, the checkbox in the Auto-Suspend section enables you to configure your system to enable other applications to access the sound system directly if it isn't being actively used. It's generally a good idea to leave this option enabled, just in case.

Selecting the Hardware tab, shown in Figure 19-10, enables you to select and configure your audio devices and sound system, as well as any MIDI devices that you're using. If you have multiple audio devices on your system, they are most easily configured using the KMix mixer application, which I'll discuss later in this chapter in the "Setting System Sound Levels" section.

FIGURE 19-10

Default hardware settings in the KDE 3 Sound settings applet



Enabling the Full Duplex option enables your sound device to both play and record sound at the same time, which can be very useful if you are using a device such an Internet phone or headset. The Custom Sampling Rate option is primarily useful for older audio devices that do not support the default sampling rate of 44100 Hz, which is CD-quality sound. However, selecting a higher sampling rate can be useful for musicians who want higher sound quality, although higher sampling rates (along with a higher setting for the Quality option) also mean increased CPU use. The Override device location option enables you to specify a device other than `/dev/dsp` (usually not necessary) and the Other custom options enables you to pass other options to the aRts daemon when it starts (also rarely necessary). After making any changes on this pane, click Apply to save them.

If you find that any of your changes to your KDE 3 sound settings are incorrect or unnecessary, you can always revert to the default KDE 3 sound settings by clicking the Defaults button on the General or Hardware tabs and clicking Apply to restore the original values.

Dealing with Multiple Audio Output Devices on KDE 3-Based Kubuntu Systems

If you have multiple audio output devices attached to your system, you may find it tricky to send output to the right device. This is especially common when you have an onboard sound card but are using a USB audio device as your primary audio output. To make sure that KDE 3 systems use the right audio device, first use the cat /proc/asound/cards command to list the audio devices that your system is aware of. Your output should look something like the following:

```
 0 [Intel ]: HDA-Intel - HDA Intel
              HDA Intel at 0x90320000 irq 22
 1 [Audio ]: USB-Audio - USB Audio
             C-Media INC. USB Audio at usb-0000:00:1a.1-1, full speed
```

This shows that my system's default device is my onboard Intel sound card, but that a USB audio output is also available. To use the USB audio output as your default, select the Override device location checkbox and identify the output that you want to use by default using the hw: keyword (or the plughw: keyword, but this doesn't always seem to work for me). For example, in the sample output shown previously, the name of my USB-Audio source is "Audio", so I would enter the string hw:Audio as my new device location. Click apply to save your changes, and your audio and video applications should begin to use the new location as their default location.

Configuring Sound on KDE 4 Kubuntu Systems

Figure 19-11 shows the default Sound configuration screen on a KDE 4-based Kubuntu system, the Device preferences screen.

As you can see, this dialog is very different that the KDE 3 version, but provides more direct control over your sound hardware and the various contexts in which sound is used. If you have multiple sound devices on your system, you can define which of them is used for Notifications, Music, Video, Communication, Games, and Accessibility by selecting the appropriate entry from the Audio Output list in the left pane. Similarly, you can specify the audio devices that are used for sound input by selecting either the Communication or Accessibility options in the Audio Capture section. After making changes to any of these, click Apply.

Click the Backend tab, shown in Figure 19-12, to modify the back end that is used by KDE 4's new Phonon sound system to handle audio and video playback.

At the moment, the Xine back end is the default, which has always worked for me, but other back ends will certainly be available in the future, so it might be handy to remember this option for the future. If you are having problems correctly displaying DVD video output, you may want to experiment with the different Deinterlacing methods that are available from the Deinterlacing method drop-down menu. The default value is to use Xine's standard video output (vo) driver, but you may find that some of the other methods work better on your system or on specific video sources. After making any changes to the settings on this tab, click Apply.

If you find that any of your changes to your KDE 4 sound settings are incorrect or unnecessary, you can always revert to the default KDE 4 sound settings by clicking the Defaults button on the Device preference or Backend tabs and clicking Apply to restore the original values.

FIGURE 19-11

Default system sound settings in the KDE 4 Sound settings applet

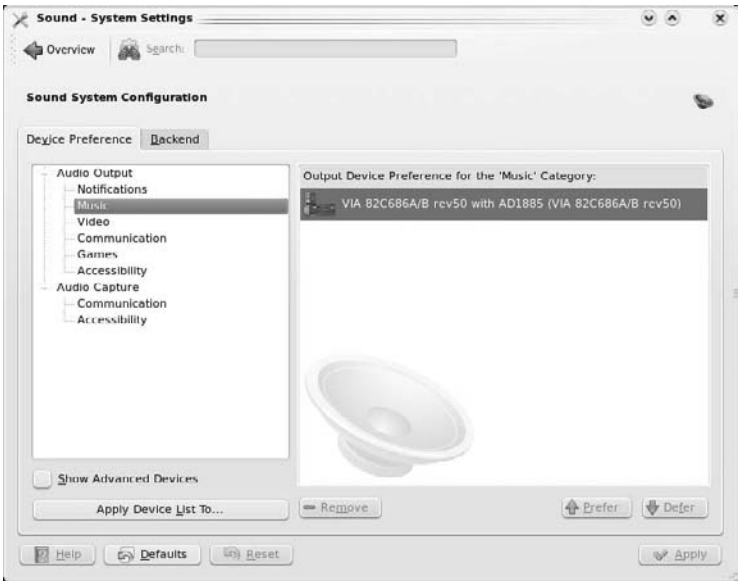
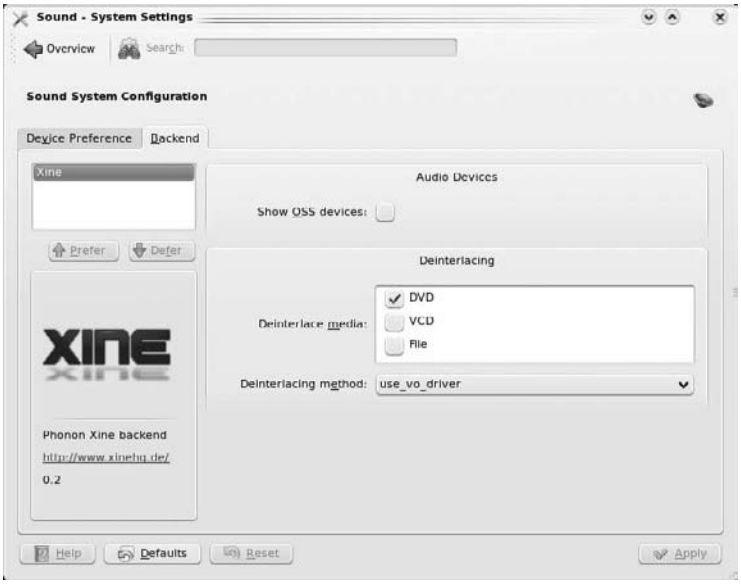


FIGURE 19-12

Back-end settings in the KDE 4 Sound settings applet



Dealing with Multiple Audio Output Devices on KDE 4–Based Kubuntu Systems

If you have multiple audio output devices attached to your system, you may find it tricky to send output to the right device by default. This is especially common when you have an onboard sound card but are using a USB audio device as your primary audio output. To make sure that KDE 4 systems use the right audio device, first use the asoundconf list command (ALSA Sound Configuration) to list the audio devices that your system is aware of. Your output should look something like the following:

```
Names of available sound cards:
Intel
Audio
```

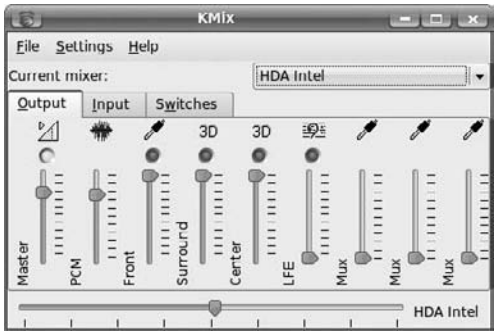
This shows that my system's default device is my onboard Intel sound card, but that a USB audio output is also available. In this example, to use the USB audio device as your default, execute the command asoundconf set-default-card Audio. Your audio and video applications should begin to use the specified sound device as the default sound device.

Setting System Sound Levels

Setting sound levels on Kubuntu systems is done through the KMix application, which looks and works almost identically on both KDE 3 and KDE 4 systems, modulo window decorations and button shapes. To start the KMix application on a KDE 3 system, select the K Menu ➪ Multimedia ➪ KMix - Mixer menu command. To start the KMix application on a KDE 4 system, select the K Menu ➪ Applications ➪ Multimedia ➪ Sound Mixer menu command. Figure 19-13 shows the KMix application on a KDE 3 system.

FIGURE 19-13

The KMix sound mixer on a KDE 3 system



If your system contains multiple sound devices, you can select the sound device that you want to control via KMix by selecting it from the drop-down list in the upper-right corner. You can then drag the sliders for any of the available tracks to adjust the sound level for that track. When you're done, you can close the KMix application—it will remain present in the System Tray section of the kicker, represented by an icon that looks like a small speaker.

To adjust sound levels from the KMix application in the kicker, left-click on the speaker icon to display a slider for the primary mixer track, in which you can drag the slider to adjust the volume. The slider pop-up also displays a Mixer button, which you can use to redisplay the full KMix application in its own window on the screen.

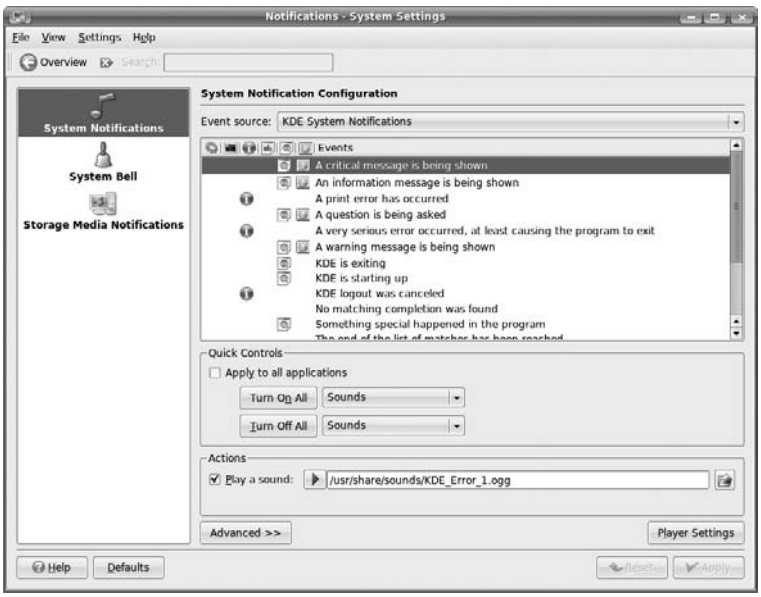
To change the default mixer track that is displayed in the KMix pop-up, right-click on the speaker icon and select the Select Master Channel menu item. This displays a pop-up listing each available mixer track for the current audio device. To change the track that is displayed in the slider, select the radio button to the left of the track that you want to use as the master volume control for the current audio device, and click OK. The newly selected track will now be displayed in the mixer pop-up that you see when left-clicking on the speaker icon.

Testing and Customizing System Sounds

You can customize the sounds and other events associated with different system and application events by selecting the System Settings ➪ Notifications icon (located in the Look & Feel section of the System Settings application. Figure 19-14 shows the configuration pane that displays after selecting this icon on a KDE 3–based Kubuntu system.

FIGURE 19-14

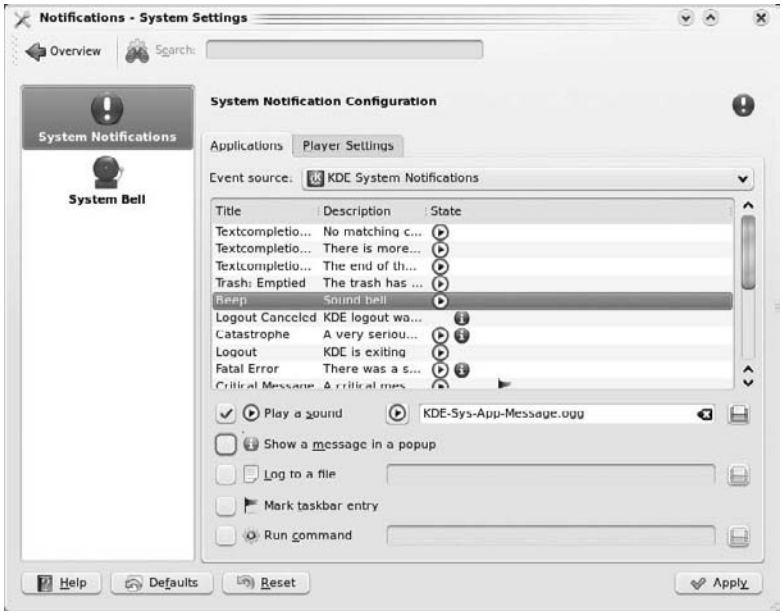
Changing system notifications on a KDE 3 system



The KDE 3 and KDE 4 versions of this configuration pane are very similar, so I'm going to focus on the KDE 3 version here. The primary difference between them (aside from graphical bells and whistles) is the fact that the KDE 3 version enables you to set different notification events and actions by clicking in the appropriate column to the left of an event, while the KDE 4 version requires that you select checkboxes for different types of notifications. These checkboxes are listed below the pane that lists available events; icons for selected notification events appear to the right of the event in the event list. Figure 19-15 shows the KDE 4 version of the Notifications configuration applet.

FIGURE 19-15

Changing system notifications on a KDE 4 system



The icons to the left of an event in the events list identify the types of notifications that are associated with each event. From left to right, these are the following:

- **Execute a program (Gear icon):** Execute a specified application when the event occurs. This can be very useful, for example, if you want to send email when a specific event occurs. (These would typically be serious system events rather than GUI events.)
- **Print to standard error output (Terminal icon):** Print a message on stderr when the event occurs.
- **Display a messagebox (Information icon):** Displays a pop-up that identifies the system event. This can be very useful for tracking network device events if you are using the console of a given machine.
- **Log to a file (Document icon):** Writes a message to a specified log file when an event occurs. This can be generally useful for auditing purposes, or if you're interested in seeing a detailed history of how you actually interact with your system. Students of Human Computer Interaction (HCI) may find this especially interesting!

- **Play a sound (Speaker icon):** Plays a specified sound when an event occurs.
- **Flash the taskbar entry (KDE icon):** Flashes the taskbar entry that is associated with an application when an event occurs. This can be very useful if the applications that you are running are minimized and an error occurs in the application or the application requires interaction.

The different classes of system and application-related events with which you can associate various forms of notification are available by clicking the Event source drop-down list and selecting the application or class of events that you're interested in. To activate a specific type of notification for a specific event, left-click in the column associated with that type of notification to the left of the event list. If you select a sound event, you can specify the sound file that is played when that event occurs by clicking the browse icon to the left of the Play a sound field and navigating to the sound that you want to play.

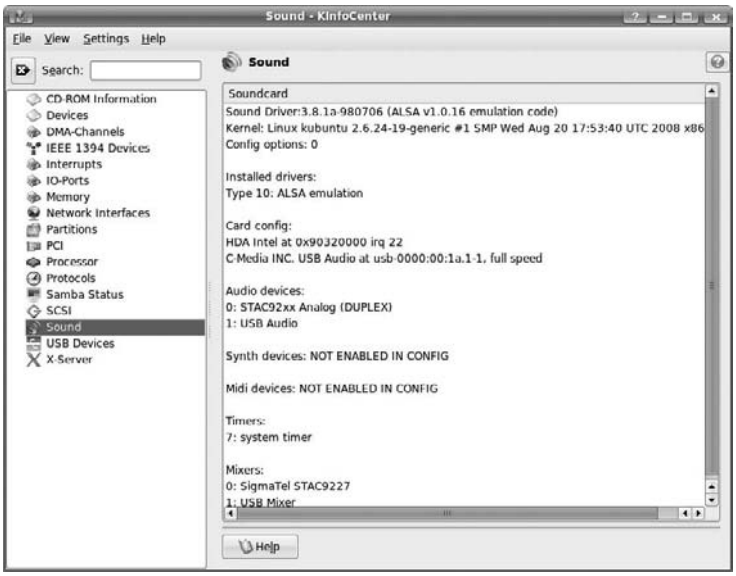
After making any changes to any of the notifications associated with various events, click Apply to save your changes. If you decide that you don't want to change things after you've already made modifications, click the Defaults button at the lower left to restore your notification settings to the system defaults.

Getting Detailed Information about Your Sound Hardware

KDE's KInfoCenter application, installed by default on KDE 3-based Kubuntu systems, provides a great overview of the devices that are present in your system and how they are configured. To start this application, select the K Menu ➪ System - KInfoCenter menu command. You can then drill down into any aspect of your system's hardware (and surprising, the status of the Samba Server) by selecting the appropriate entry from the left pane. Figure 19-16 shows the KInfoCenter application displaying detailed information about the sound devices, drivers, and related subsystems on one of my Kubuntu systems.

FIGURE 19-16

Hardware information on a KDE 3–based Kubuntu system

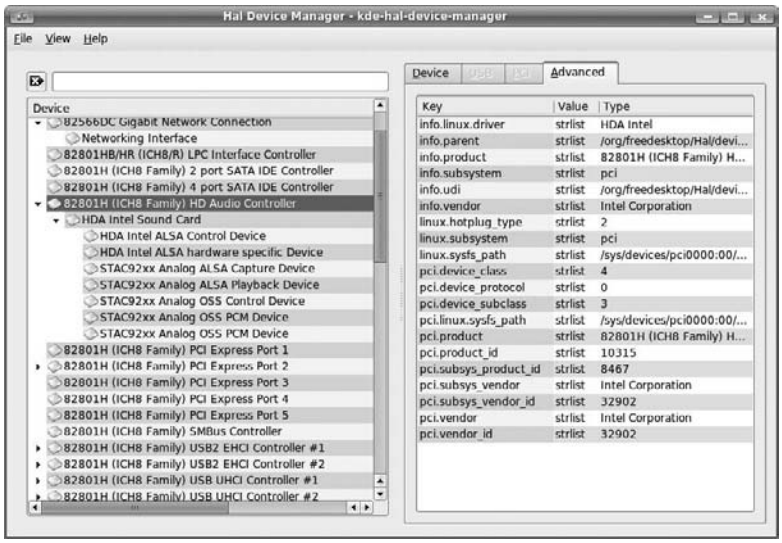


This application is also available for KDE 4-based Kubuntu systems, although it is not installed by default. You can install it by using any of the applications discussed in Chapter 23 to install the kinfocenter package.

You can obtain much more detailed information about your system's hardware and its configuration by installing the `kde-hal-device-manager` package, which will create a K Menu ➪ System ➪ Device Manager entry. (This menu entry is called Hardware Viewing on KDE 4–based Kubuntu systems.) The HAL (Hardware Access Layer) Device Manager utility is available for both KDE 3 and KDE 4, but it is not installed by default on Kubuntu systems. To install this package, use any of the software utilities discussed in Chapter 23. Figure 19-17 shows the HAL Device Manager displaying detailed information about sound hardware (on a different system than that shown in Figure 19-16).

FIGURE 19-17

Detailed hardware information on a KDE 3–based Kubuntu system



The illustrations in this section show the KDE 3 versions of these utilities, but the KDE 4 versions are effectively identical, modulo window decorations and themes.

Setting Blank CD and DVD Preferences

Now that few people bother with floppy disks (or, perhaps, know what they are), CDs and DVDs are the most common examples of removable media that you'll use with your computer system. Removable media is the generic term for any media that contains digital data and that is not a mandatory part of your computer system. Removable disks, flash drives, and other devices that show up as disk drives (like an iPod) are not exactly the same things as far as Ubuntu is concerned and are discussed in Chapter 21, "Consumer Electronics and Ubuntu," as well as in Chapter 24, "Adding Hardware and Attaching Peripherals." This section focuses on how you can configure what your Kubuntu system does when you insert a prerecorded or blank CD or DVD.

NOTE **This section discusses what happens on a KDE 3-based Kubuntu system. At the moment, KDE 4-based Kubuntu systems don't automatically offer to start any applications for you when you insert a blank CD or DVD—you'll have to start Amarok, k3b, Dragon, or kscd manually by selecting the appropriate entry from the K Menu ➪ Applications ➪ Multimedia menu.**

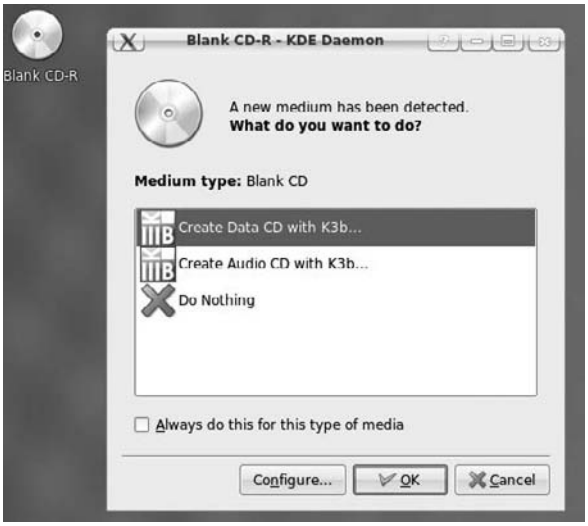
When you insert a blank CD or DVD in your Kubuntu system, the system does two things:

- Displays an appropriate icon for the type of blank media that you have inserted
- Displays a dialog asking what you would like to do with that media

Figure 19-18 shows the icon and dialog displayed when you insert a blank CD-R or CD-RW disc in a compatible drive on a KDE 3–based Kubuntu system. A very similar icon and dialog are displayed when you insert a blank DVD-R disc into a compatible drive on your system.

FIGURE 19-18

Kubuntu Dialog and icon for blank CD-R/RW insertion

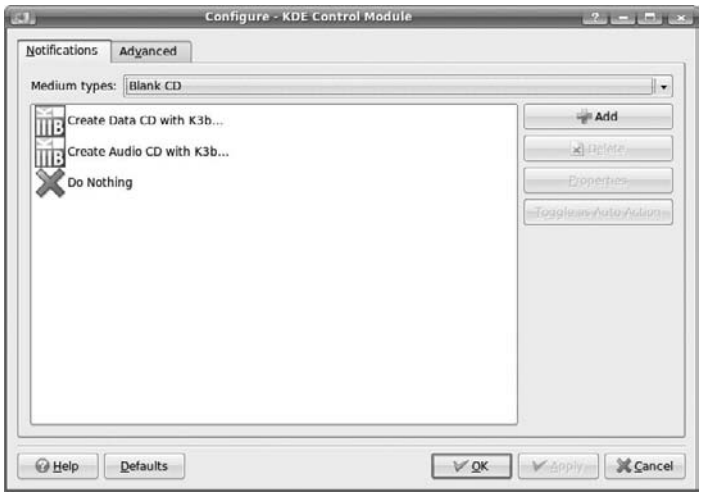


In either case, clicking Cancel closes the dialog without starting an application. You can select either the Create Data CD with k3b or Create Audio CD with k3b checkboxes (or their DVD equivalents) to start the excellent k3b application (discussed later in this chapter) to begin defining the contents of your CD or DVD in either format. You can also click the Always do this for this type of media checkbox to automatically start k3b in one of those modes whenever you insert a blank disc. You can also select Do Nothing and click OK to proceed if you have something else in mind for your blank media or simply don't want to start k3b at this time.

To prevent your system from displaying the dialogs shown in Figure 19-8, select the Do Nothing entry and click the Always do this for this type of media checkbox. To add items to this menu, for example if you want to use a different application to work with Blank CDs, click Configure to start the KDE Control Center's Configure utility (KDE 3–based systems only), shown in Figure 19-19.

FIGURE 19-19

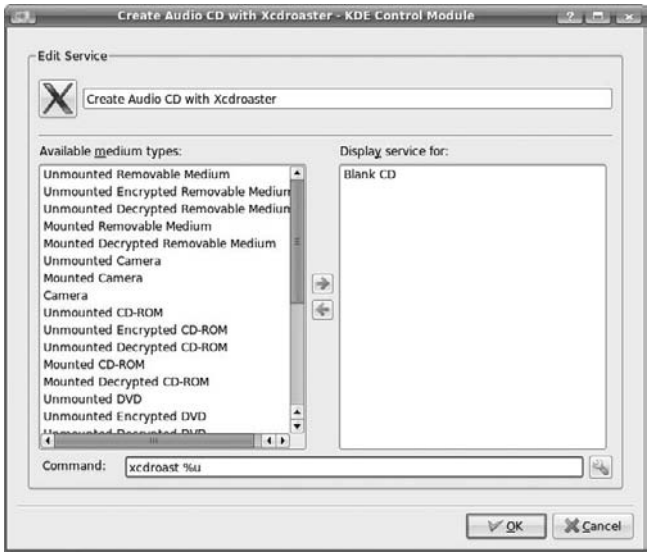
Setting Kubuntu preferences for blank CDs or other media



Once this dialog displays, select the entry that you want to modify, either Blank CD or Blank DVD from the Media Type drop-down menu, and click Add to define a new action. A dialog like the one shown in Figure 19-20 displays.

FIGURE 19-20

Creating a new action for blank CDs on Kubuntu systems



To define a new action for a Blank CD, double-click on Blank CD from the scrollable Available medium types list at the left of this dialog. You can then specify the menu entry that you want to display at the top of this dialog, click the icon at the upper left to select a custom icon, and specify the command line that you want to use in the Command field at the bottom of this dialog. As an example (not a recommendation), Figure 19-20 shows an entry being created for the venerable CDroaster application, an ancient but capable CD-burning utility. After defining a new entry, click OK to save it, and click OK again to exit the Configure utility. If you had followed along with this example, you would see the dialog shown in Figure 19-21 the next time that you inserted a blank CD.

FIGURE 19-21

The updated Kubuntu dialog that appears when a blank CD is inserted



Unfortunately, configuring default actions for removable media events is one of the areas in which KDE 4 still seems to need some work. However, given its glitzy look and feel and the general friendliness of KDE, I'm sure that this is only a matter of time (and programming effort, of course).

For information about playing or burning CDs and DVDs, see the sections later in this chapter entitled "Working with CDs" and "Working with DVDs."

Exploring the Linux Multimedia Stack

Like related concepts such as disk access and displaying graphical output, your system must be able to play multiple audio steams at the same time, even for something as simple as playing the sound associated with a system event while also playing your favorite CD. The single-tasking approach, where ancient audio applications seized sole control of an audio output device, simply isn't practical today. For this reason, most audio playback on Linux systems is handled by an intermediate layer or sound server, which can play and mix multiple audio streams at the same time.

Introducing Drivers, Servers, and Frameworks

I doubt that it will come as a surprise that Linux distributions have taken a number of different approaches to supporting audio playback over the years. Hacking through the morass of acronyms that seems especially deep in the Linux multimedia space, the next few sections explain the core concepts in Linux Multimedia.

Linux Sound APIs

The primary sound APIs used on Linux systems (and UNIX systems, if anyone still cares) are OSS (Open Sound System) and ALSA (Advanced Linux Sound Architecture). These are the lower-level drivers that enable communication with sound devices. As the name suggests, ALSA is a Linux-only API, while OSS can be used on most UNIX-like systems.

OSS was the standard sound API on Linux systems through the 2.4 Linux kernel series, after which ALSA became the standard sound API. ALSA includes an OSS emulation mode for backward compatibility.

OSS is still being developed, although it has been marked as deprecated in the 2.5 and newer Linux kernels. It remains a more portable API if you are writing code that is intended to run on multiple UNIX-like systems, not just Linux systems.

Linux Sound Servers

Linux systems use sound servers (AKA sound daemons) to enable the simultaneous playback of multiple output devices and to enable playing and controlling sound over a network. (ALSA audio drivers inherently support simultaneous playback and mixing, but most OSS drivers do not.)

Traditionally, the most common sound server on GNOME-based systems has been ESD (Enlightened Sound Daemon, aka EsounD, `www.tux.org/~ricdude/overview.html`). With Ubuntu 8.04 and in most other up-to-date GNOME-based Linux distributions, ESD has been replaced by PulseAudio (`www.pulseaudio.org`), though the transition was not necessarily a smooth one. Don't worry—subsequent sections discuss PulseAudio in more detail, though in my experience "it just works" on Ubuntu 8.10 (Intrepid) systems.

On KDE-based systems such as Kubuntu, the story is, of course, slightly different. KDE 2 and 3 systems used aRts (analog real-time synthesizer sound server, `www.arts-project.org`) on which development has now halted. KDE 4 and better systems will use the new Phonon Multimedia API (`phonon.kde.org`). Whether KDE 4 will adopt PulseAudio and use it by default with back-end frameworks such as GStreamer and Xine is still up in the air, although a preliminary version of PulseAudio for the Xine back end is already available.

Just to keep things confusing and infinitely flexible, several other, system-independent sound servers such as the JACK Audio Connection Kit (`www.jackaudio.org`) and NAS (Network Audio System, home pages) are also available and have been around for a while.

Linux Multimedia Frameworks

GStreamer (`www.gstreamer.net`) is a pipeline-based multimedia framework that can be used by applications to process the data between when it is read from ALSA drivers and before it is output, either directly or through a sound server. GStreamer uses a plug-in-based architecture that makes it easy to extend GStreamer to support new audio or video formats. Though originally developed for use on GNOME-based systems, GStreamer was designed to be platform-independent and is thus used on platforms that include Linux, Mac OS X, Microsoft Windows, and Solaris. The GStreamer framework is also being used by the KDE 4 Phonon project.

Some multimedia applications provide their own multimedia frameworks, which are therefore highly portable to other platforms and which can often be accessed by other applications. Good examples of this are the MPlayer Movie Player (`www.mplayerhq.hu`), VLC Media Player (`www.videolan.org/vlc`), and the Xine Multimedia player (`http://xinehq.de`). For example, the current Phonon implementation on KDE 4–based Kubuntu 8.10 systems uses the Xine back end for audio and video output.

What Should You Care About?

Having learned all of these acronyms, how do they affect you?

- If you simply want to play audio on your Ubuntu or Kubuntu system, all you really need is ALSA.
- Ubuntu users who want to play MP3 files or watch DVDs, or who might ever encounter an audio or video format that is not completely open source should install GStreamer plug-ins to expand the core GStreamer capabilities that are present on your system.
- If you need networked sound or want to be able to separately control the audio levels of different audio applications, you should take advantage of PulseAudio on Ubuntu systems.
- If you're using a KDE 3-based Kubuntu system (Kubuntu 8.04), PulseAudio doesn't apply.
- If you're using a KDE 4-based Kubuntu system (Kubuntu 8.10), PulseAudio support for the Phonon back ends wasn't ready for prime time when this book was written. I'd suggest searching the Kubuntu site and the rest of the Web for more information before experimenting with PulseAudio on a KDE 4-based Kubuntu system.

Dealing with Proprietary Audio and Video Formats

The rest of this chapter provides detailed information about installing and configuring various packages and applications in order to enable you to work with various types of audio and video on your Ubuntu and Kubuntu systems. This includes MP3 files that you've purchased, DVDs that you've purchased, and so on. You'd think that this would "just work"—but no! The music and motion picture industries and some bogons who have copyrighted various file formats impose many hurdles for the innocent Linux consumer to overcome in order to use legally purchased products. In order to be able to play many types of audio files, work with DVDs, and so on, you must install software packages on your Ubuntu or Kubuntu system that the folks at Canonical cannot install by default for legal reasons. You can install these packages by installing the *ubuntu-restricted-extras* package on Ubuntu systems, or the *kubuntu-restricted-extras* package on Kubuntu systems. These packages include a base set of GStreamer plug-ins (see the next section, "Configuring the GStreamer Framework and Plug-ins," for more information about GStreamer) that will enable you to play most unencrypted audio and video formats.

Please note that I am not a lawyer and that installing the software that you need to use in order to play legitimately purchased audio and video data is completely your responsibility, regardless of how silly that may be. Installing these packages on your system may someday cause jackbooted enforcers from the RIAA and MPAA to kick down your door and drag you off to be forcibly reeducated. Hordes of winged monkeys with legal degrees may darken the skies over your home forever. You should probably never install the `restricted-extras` packages on a business system. For information about legally playing various proprietary audio formats on Ubuntu 8.04 and 8.10 systems, see the "Getting Commercial GStreamer Plug-Ins" sidebar at the end of the next section.

Configuring the GStreamer Framework and Plug-Ins

As introduced in the previous section, most GNOME audio and video applications use an underlying open source multimedia framework called GStreamer (`http://gstreamer.freedesktop.org` or `www.gstreamer.net`). GStreamer makes it easy to integrate codecs by writing simple plug-ins that integrate with its framework. Most of the GNOME applications discussed in this chapter leverage the GStreamer framework for codec support via plug-ins. For this reason, you should use your favorite package management software (discussed in Chapter 23) to install the latest versions of GStreamer and associated sets of plug-ins before starting Rhythmbox, trying to play MP3 files, attempting to convert sound files, or attempting to play a DVD. If you don't, you may be extremely disappointed in Ubuntu's audio and video capabilities—at least, until you install the GStreamer plug-ins and try again.

As explained in Chapter 1, the Ubuntu folks are very strict about the legal and licensing requirements of the packages that they include in the default Ubuntu distribution. Because of copyright and licensing issues, many of the codecs that are commonly used by audio applications on other platforms are not installed by default on your Ubuntu Linux system. The best example of this is anything having to do with the MP3 audio format. Not everyone needs to play the most common audio files format found in the known universe, right? I can't blame them for this perspective, but in this case, it doesn't meet my needs.

To make GStreamer as full-featured as possible, use your favorite package management software from Chapter 23 to install the following packages (or to make sure that they're already installed):

- `gstreamer0.10-ffmpeg`
- `gstreamer0.10-plugins-bad`
- `gstreamer0.10-plugins-bad-multiverse`
- `gstreamer0.10-plugins-base`
- `gstreamer0.10-plugins-base-apps`
- `gstreamer0.10-plugins-good`
- `gstreamer0.10-plugins-ugly`
- `gstreamer0.10-plugins-ugly-multiverse`
- `gstreamer0.10-tools`
- `totem-gstreamer`

TIP **You may not need all of these unless you are converting audio files from one format to another, and even in that case, the plug-ins that you need depend on the types of files that you're converting. I'm assuming that you would rather err on the side of caution—that is, installing a version of the GStreamer framework that is as full-featured as possible, at the expense of a bit of disk space.**

GStreamer is a library that is used by many applications, not an application itself. The GStreamer library enables applications to access GStreamer plug-ins. Having installed the suggested GStreamer plug-ins, applications such as Rhythmbox and SoundConverter, discussed later in this chapter, will now be able to work with most types of audio files.

TIP **For up-to-date information on using GStreamer-based applications to play proprietary audio and video formats, see** `https://help.ubuntu.com/community/RestrictedFormats`.

Getting Commercial GStreamer Plug-ins

In September, 2008, Canonical began to offer commercial 32-bit codecs (actually just DECs, since they are only decoders) for various Microsoft Windows and Motion Picture Expert Group (MPEG) media formats through its online store (`http://shop.canonical.com`; 64-bit versions are coming RSN). These codecs, implemented as GStreamer plug-ins and thus intended for use on Ubuntu systems (not Kubuntu systems), are from Fluendo (`www.fluendo.com`), a Linux and UNIX multimedia firm that has already developed and released an open source MP3 codec (`www.fluendo.com/press/releases/PR-2005-05.html`). Two commercial packages are available:

- **Fluendo Windows Media and MP3 Playback Pack**: Provides codecs for Windows Media Audio (WMA files for Windows Media 7, 8, 9, 10, Pro, Lossless, and Speech), Windows Media Video (WMV files for Windows Media 7, 8, 9, and VC1), Windows Media MMS, Windows Media ASF, and MP3 Audio.
- **Fluendo Complete Playback Pack**: Provides codecs for Windows Media Audio (WMA files for Windows Media 7, 8, 9, 10, Pro, Lossless, and Speech), Windows Media Video (WMV files for Windows Media 7, 8, 9, and VC1), Windows Media ASF, Windows Media MMS, MPEG2 Video, MPEG4 ISO and Part 2 (DivX version 4, 5, and 6) Video, H.264/AVC Video, MP3 Audio, and AAC Audio.

Needless to say, this caused the open source community to burst into flames, with pundits, open source purists, and Ubuntu fans turning on each other like members of the Donner party. Free GStreamer plug-ins that support most of these formats are already available but (as mentioned in the previous sections) cannot be used legally in many countries.

My take on this? Get over it! People who want to be able to replace Microsoft Windows installations with Ubuntu installations in commercial or off-the-shelf environments need these codecs. Just because no one likes the fact that certain audio and video formats are proprietary doesn't mean that one can ignore that fact in enterprise deployments. Similarly, no one is going to buy a netbook or mobile Internet device that can't play MP3s. Most of us don't have to buy these codecs, but having legally licensed versions available will spur the adoption of Ubuntu in business, academic, and preinstalled, off-the-shelf environments. That's a good thing for all of us.

Configuring PulseAudio on Ubuntu

PulseAudio, introduced in the beginning of this section, is the latest sound server that is available for Ubuntu and most other GNOME-based Linux distributions. Originally called Polypaudio, the name was changed because Polypaudio sounded too much like something that you'd want to have surgically removed.

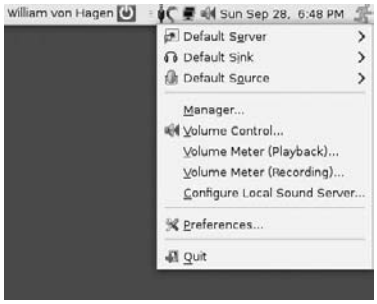
The PulseAudio server, a GStreamer plug-in for PulseAudio integration, and a startup item for integration between X Window system events and PulseAudio are all installed by default on both Ubuntu 8.04 (Hardy) and 8.10 (Intrepid), but PulseAudio only works out of the box on Intrepid systems. (See the next section for the details of getting PulseAudio working on Ubuntu Hardy systems.)

On both Ubuntu Hardy and Intrepid systems, the default PulseAudio installation provides a number of command-line utilities for configuring and monitoring PulseAudio. However, if you're like me, you'll probably want to install the `padevchooser` package, which provides graphical utilities such as the PulseAudio Device Chooser (`padevchooser`), PulseAudio Manager (`paman`), PulseAudio Preferences (`paprefs`), PulseAudio Volume Control (`pavucontrol`), and PulseAudio Volume Unit Meter (`pavumeter`)

applications. Installing the `padevchooser` package using your favorite package management utility (discussed in Chapter 23) creates the Applications ➪ Sound & Video ➪ PulseAudio Device Chooser menu item, which starts the PulseAudio Device Chooser as an applet in the notification area portion of the top panel. Its icon looks like an audio jack and a piece of audio cable. Left-clicking this icon displays the menu shown in Figure 19-22.

FIGURE 19-22

The GNOME panel's PulseAudio Device Chooser menu

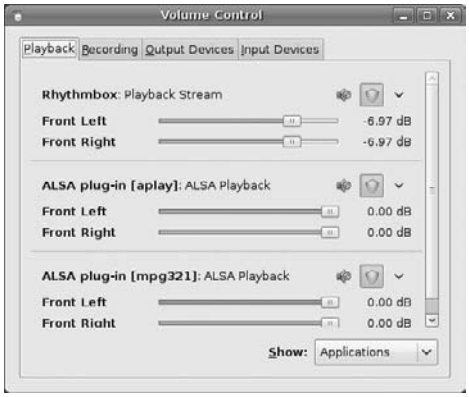


The items on this menu provide access to the graphical PulseAudio utilities that were installed as part of the `padevchooser` package. For example, selecting the Volume Control menu item displays the PulseAudio Volume Control application, which provides a Playback tab where you can set the volume for any audio device that is currently playing audio. Figure 19-23 shows this tab when three different applications are all playing through PulseAudio at the same time.

To adjust the volume of any of these, simply drag the slider for that application. You can also click the speaker icon beside the sliders for any application to temporarily mute its audio stream. Very cool.

FIGURE 19-23

Separately adjusting the volume for audio applications in PulseAudio



If you decide that you like the PulseAudio Device Chooser application (and what Ubuntu PulseAudio user wouldn't?), you can configure it to always start in any Ubuntu login session by selecting the PulseAudio Device Chooser's Preferences menu item to display the PulseAudio preferences application, shown in Figure 19-24.

FIGURE 19-24

Setting PulseAudio preferences



Select the Start applet on session login checkbox to automatically start the PulseAudio Device Chooser each time that you log in on your Ubuntu system.

If you're still running Ubuntu Hardy (8.04), see the sidebar entitled "PulseAudio on Ubuntu 8.04 Systems" for some tips and tricks that you'll need to use there.

PulseAudio on Ubuntu 8.04 Systems

PulseAudio is installed by default on Ubuntu 8.04 systems, and almost works correctly out of the box. I've always had to make a few tweaks to get it to run correctly. The most important of these is to configure ALSA to work through the PulseAudio server by executing the following command:

```
asoundconf  set-pulseaudio
```

This creates two startup files in your home directory that ALSA will use when it starts:

- **.asoundrc:** ALSA's primary startup file. Unless you add anything to the file, this file simply invokes the other startup file that was created by the `asoundconf` utility.
- **.asoundrc.asoundconf:** The configuration file that tells ALSA to use PulseAudio via the following two lines:

```
pcm.!default { type pulse }
ctl.!default { type pulse }
```

After running this command, you can restart ALSA using the command `sudo /etc/init.d/alsa-utils restart`). I also generally log out and log back in at this point to make sure that ALSA picks up the new settings automatically upon login.

After configuring ALSA to use PulseAudio, you may see errors when you try to use ALSA-specific utilities such as `amixer`, the ALSA mixer application. To resolve this, use your favorite package management utility to install the `asound2-plugins` package, which provides various libraries for integrating ALSA with sound servers. Similarly, I also usually install the `asoundconf-gtk` application, which creates the System ➪ Preferences ➪ Default Sound Card menu item that starts a utility which enables you to quickly switch between PulseAudio and ALSA. Although PulseAudio is a great sound server, you may not always want or need to run one. For example, I run PulseAudio on my laptop when it's docked, but I only use ALSA when I'm on the road. No need to run an extra process when I'm the only possible user.

For additional information about PulseAudio configuration on Ubuntu Hardy systems, see the Perfect Setup topic on the PulseAudio wiki (`www.pulseaudio.org/wiki/PerfectSetup`).

Playing CDs and DVDs

I still have a few friends who use Windows voluntarily, and one of their favorite rationales is the statement, "It just works!" Thanks to the efforts of all the contributors and testers who are associated with the Ubuntu and Kubuntu projects, playing music on those systems has this same mantra—I have never had a problem with an internal CD drive on any desktop or laptop system where I have installed Ubuntu or Kubuntu Linux.

Unfortunately, playing DVDs is a little trickier. The only thing that could possibly make anyone nostalgic for videotapes is how badly the entertainment media is trying to screw DVD users. Region-specific DVDs. Video encryption. Licensed codecs. Suing young adults and prohibiting free speech for academicians just because both are smarter than their research staff. Need I say more?

For all of these reasons (and more), playing DVDs on a computer system isn't always an easy thing to do. However, Ubuntu provides a great DVD player called Totem that provides high-quality playback capabilities. Totem can either leverage the GStreamer framework that I discussed earlier in this chapter or leverage other multimedia back ends such as Xine. Similarly, KDE 3-based Kubuntu systems provide the Kaffeine media player, and KDE 4-based Kubuntu systems provide Dragon, both of which are equally capable.

Configuring how your system responds when you insert an audio or data CD, a prerecorded DVD, or a data DVD was discussed earlier in this chapter in the Ubuntu or Kubuntu sections entitled "Setting Blank CD and DVD Preferences." To summarize, Ubuntu and KDE 3-based Kubuntu systems display a pop-up dialog that enables you to select the application that you want to use to play your CD or DVD. KDE 4-based Kubuntu systems didn't display this menu at the time that this book was written, but I'd expect them to eventually.

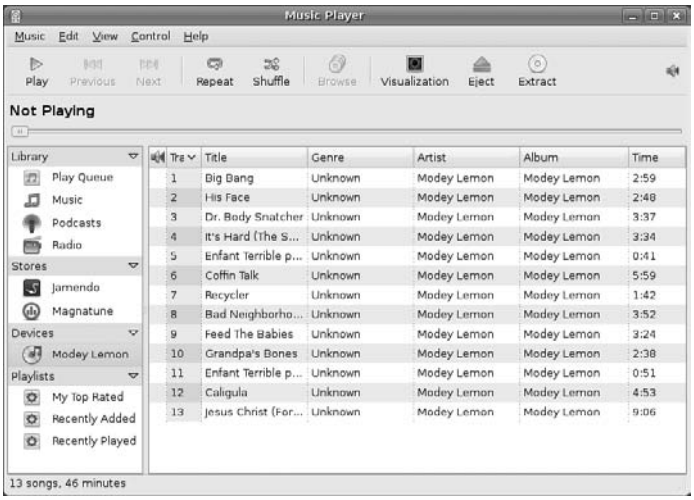
- If no pop-up dialog displays when you insert a CD or DVD, you can select the application that you want to start from Ubuntu's Applications ➪ Sound and Video menu, the K Menu ➪ MultiMedia menu on KDE 3-based systems, or the K Menu ➪ Applications ➪ MultiMedia menu on KDE 4-based Kubuntu systems.

Playing Audio CDs Using Rhythmbox on Ubuntu

Once Rhythmbox starts, you may need to select the entry in the Device section in the pane at the left to display information about the CD that you've inserted. This will initially try to look up any CD that you insert in the free online CD information database (freedb.org). If it finds a matching entry, it displays information about the CD and the tracks that it contains, as shown in Figure 19-25.

FIGURE 19-25

Information about a CD from FreeDB



If the CD is not listed in FreeDB, Rhythmbox displays a similar screen, using a generic entry for the CD that says *Unknown Audio* and *Unknown* for both the Artist and Title in its track listing window, assigning each track a generic name such as *Track 1*, *Track 2*, and so on.

If you simply want to play the CD, click the Play icon in the toolbar to play the CD. It just works! If your CD was found in the FreeDB database, Rhythmbox will also display a miniature version of the CDs cover art in the lower-left corner of the Rhythmbox window. As you might expect, you can play specific tracks by double-click on the track that you want to play.

Once you're done playing your CD, you can click the Eject icon in the toolbar to eject this CD and insert another. To exit Rhythmbox, select the Music ➪ Quit menu command. If you want to close the Rhythmbox window but continue to play the CD, you can either minimize the window, using the standard window controls, or select the Music ➪ Close menu item. In either case, the Rhythmbox icon (which looks like a speaker) will still be displayed in the system tray portion of the GNOME panel.

Playing Audio CDs with Amarok on Kubuntu

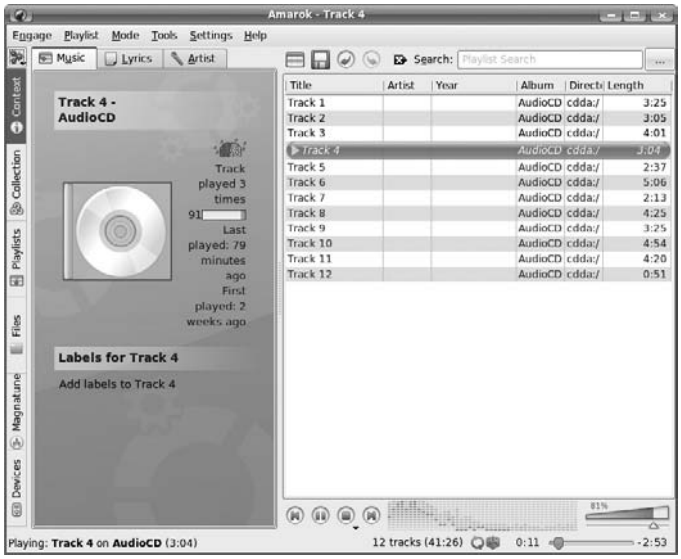
Once Amarok starts, it displays its startup screen and tries to look up your CD in the free online CD information database freedb.org, or at least claims to do so. If it finds a matching entry, it displays information about the CD and the tracks that it contains. However, I've never had much luck getting Amarok

to actually retrieve this information for me on the latest KDE 3–based Kubuntu systems, although it does a good job of this on KDE 4–based Kubuntu systems. I typically don't worry about this, because I usually just rip CDs for convenience sake when I first get them (using k3b, as discussed in the section of this chapter entitled "Ripping Audio CDs using k3b on Kubuntu"). As I'll discuss the section entitled "Playing Audio Files and Internet Radio on Kubuntu," Amarok does a great job of dealing with most types of online audio—it just seems to have some issues with the titles of audio CDs.

If Amarok manages to find information about your CD, it will display song titles in its main window. Otherwise, it will display a generic entry for the CD that assigns each track a generic name such as *Track 1*, *Track 2*, and so on, as shown in Figure 19-26.

FIGURE 19-26

Generic audio CD information in Amarok



Playing CDs is quite straightforward. Amarok will start playing the CD immediately after trying to do the freedb.org lookup. You can use the controls below the playlist window to pause or stop playing your CD, and can use the items on the Engage menu to move between tracks as well as start and stop playing the CD. You can also click on the name of a track to begin playing it, but this is not very useful if you don't know the real names of the tracks you're listening to.

NOTE **If you want to close the Amarok window but continue to play the CD, you can use the standard window control to close the window. The Amarok icon (which looks like a wolf's head) will still be displayed in the system tray portion of the KDE panel.**

Once you're done playing a CD, you can eject the disc manually and insert another. To exit Amarok entirely at any point, select the Engage ➪ Quit menu command or right-click on the Amarok icon in the system tray and select Quit from the content-sensitive menu.

If you really want to play audio CDs on KDE 3–based systems, I'd suggest using Kaffeine instead, since it is much better at actually communicating with freedb.org. After starting Kaffeine from the K Menu ➪ Multimedia menu, click the Audio CD button to retrieve information about your CD and begin playing it. Figure 19-27 shows a CD being played in Kaffeine, complete with artist, title, and track names.

FIGURE 19-27

Playing an audio CD in Kaffeine



Installing Software to Play Encrypted DVDs

Installing the library necessary to read encrypted DVDs gives you the power to actually watch the DVDs that you've legally purchased, so let's install it now.

CAUTION **In my personal opinion, installing the CSS (`http://en.wikipedia.org/wiki/Content_Scramble_System`) software that enables you to play commercial DVDs on your Ubuntu Linux system is a right, not a privilege or a crime. However, I am not a lawyer. Installing this software on your system may someday cause jackbooted enforcers from the MPAA to kick down your door and drag you off to be forcibly reeducated. You should make sure that you are comfortable with installing this software, and should probably never install it on a business system.**

To install this library, execute the following command from a GNOME terminal, KDE Konsole, or xterm command line:

```
$ sudo /usr/share/doc/libdvdread3/examples/install-css.sh
```

As always, you'll be prompted for your password by the sudo command, which will then execute the specified shell script. The install-css.sh script will both retrieve the decss library and install it on your system for you.

Congratulations! You can now watch your favorite DVDs in the privacy of your personal computer system.

Playing DVDs with Totem on Ubuntu

The GNOME Totem application (`http://hadess.net/totem.php3`) is a full-featured DVD and video file player that relies on underlying technologies such a GStreamer (the default) or Xine (another video display and codec engine) for many of its low-level capabilities. Totem is preinstalled on Ubuntu Linux systems, as it is on all GNOME-based Linux distributions.

Using Totem to play a DVD is easy. Because most prerecorded DVDs feature some sort of encryption, you should install the GStreamer framework and associated libraries before trying to play a DVD. For details on how to do this, see "Installing the GStreamer Framework and Plug-Ins," earlier in this chapter.

As explained earlier in this chapter, Ubuntu's GNOME desktop comes preconfigured to start Totem for you when it detects that you have inserted a prerecorded DVD. Therefore, if you haven't changed any of those settings, the only remaining step in playing a prerecorded DVD is the following: Insert it into your DVD drive. GNOME starts Totem automatically for you, and you'll see a screen like that shown in Figure 19-28. (The DVD that is displayed will probably be different, of course.)

TIP **You can also start Totem manually at any time by selecting the Applications ➪ Sound & Video ➪ Movie Player menu command, or by typing *totem* in a GNOME Terminal or xterm terminal session.**

FIGURE 19-28

A prerecorded DVD playing in Totem



Like most GNOME applications, Totem supports a variety of keyboard commands that make it easy for anyone to use. Some of the more common keyboard commands in Totem are the following:

- **B:** Skip back to the previous chapter or movie on the DVD.
- **F:** Switch to full-screen mode. In full-screen mode, Totem displays an icon in the upper-right corner of the screen that you can click on to return to the original display mode shown in Figure 19-28.
- **H:** Hide Totem's controls. Pressing H again redisplays them.
- **N:** Skip to the next chapter or movie on the DVD.
- **F9:** Hide Totem's contents sidebar. Pressing F9 again redisplays the sidebar.
- **Page-Up:** Increase the volume.
- **Page-Down:** Decrease the volume.

Once you have watched a movie in Totem, you can select the Movie ➪ Eject menu command (or press Ctrl+E) to eject the current disk. If you choose to watch another, you can select the Movie ➪ DVD command to open and begin playing that DVD. If you do not want to watch another, you can exit from Totem by selecting the Movie ➪ Quit menu command.

TIP **Totem is still actively under development, and some versions have problems opening certain prerecorded DVDs when using the menu commands. If Totem complains about not being able to locate the codecs required to play a DVD in response to the Movie ➪ DVD command, exit Totem, eject the DVD, and let GNOME start Totem correctly for you when you reinsert the DVD.**

Playing DVDs on Kubuntu Systems

The change from KDE 3 to KDE 4 is much more than just a new level of eye candy. The new graphical and multimedia frameworks that have been introduced with KDE 4 have brought with them many new applications, which were designed from the ground up with these new frameworks in mind. At the same time, many KDE 3 applications are moving to KDE 4, but that migration takes time because KDE 3 and 4 are so different internally, especially for multimedia applications.

The applications used to play DVDs on KDE 3 and KDE 4 systems provide an excellent example of this. Kaffeine is an excellent media player that has been a favorite of KDE 3 fans for quite a while, but its move to KDE 4 is taking some time. This temporary gap has provided an opportunity for Dragon, another excellent DVD and video app, to get its chance at stardom. One of the greatest things about Linux and open source software in general is its diversity and the sheer number of applications of every sort that are available. This can seem overwhelming at first, I like to think of it as being a permanent guest at the world's greatest free smorgasbord, where you get to see and try things you might not have even heard of before, and can still always fall back to your favorites when they finally make it to the buffet table.

Playing DVDs Using Kaffeine on KDE 3-Based Systems

Kaffeine (`http://kaffeine.kde.org`) has been a multimedia star for several KDE releases now, serving as the standard media player on most KDE-based Linux distributions, although excellent media players such as the audio-only JuK (`http://developer.kde.org/~wheeler/juk.html`) and video-only Noatun (`http://noatun.kde.org`) players are also quite popular.

Inserting a DVD in your drive on a KDE 3-based Kubuntu system and selecting Play DVD with Kaffeine starts Kaffeine in video mode. The startup screen that you see when you start Kaffeine from the K

Menu ➪ Multimedia menu on KDE 3-based Kubuntu systems (shown long ago in Figure 19-27) doesn't appear in this case—Kaffeine automatically begins playing your DVD—usually. If you see the warning shown in Figure 19-29, you didn't follow the instructions given in the section entitled "Installing Software to Play Encrypted DVDs" earlier in this chapter, and you can either follow those instructions or the instructions in the dialog shown in Figure 9-29. Funny how it's necessary to install a special package to enable you to play legally purchased DVDs on a legally purchased computer, but that's the MPAA and corporate greed for you (and is not the fault of Linux, Kaffeine, or anything else with a frontal lobe).

FIGURE 19-29

Warning! It may be illegal to play legally purchased DVDs!



After the `css` (Content Scrambling System) package is installed, playing a DVD with Kaffeine will look like the screen shown in Figure 19-30.

FIGURE 19-30

Playing a DVD in Kaffeine



Like most KDE applications, Kaffeine is quite easy to use. Kaffeine's Player ➪ DVD menu provides submenus that enable you to jump to specific titles or chapters in your DVD. The Ctrl+Shift+F keyboard combination enables you to switch to full-screen mode and back again. Pressing the M key enables you to hide or display the Kaffeine menu bar and tabs.

Like many video players, Kaffeine uses the Xine display engine, which uses a hardware-accelerated video driver to play DVDs and videos at the right speed. One interesting side effect of this is that if you ever attempt to capture a screenshot of Kaffeine or any of the DVDs you're watching, all you'll see where the video should be is a blue screen because the video is playing at a different level. If you ever need to do a screen capture of a video, you can change the video driver that Xine uses by selecting the Settings ➪ xine Engine parameters menu item, selecting the video icon in the left-hand toolbar, and selecting the xshm driver from the drop-down menu. After clicking Apply (or OK to close the preferences dialog), you'll be able to take a screenshot of your favorite scene from your favorite movie.

Playing DVDs Using Dragon on KDE 4-Based Systems

Dragon Player (`http://dragonplayer.org`) is a simple DVD and video media player that was designed for KDE 4, and is therefore the default DVD/video player on Kubuntu 8.10 systems. Dragon uses the Xine video multimedia back end that is current the primary video back end for KDE 4, so you should make sure that you install the `kubuntu-restricted-essentials` and `libxine1-all-plugins` packages before attempting to play a DVD on your KDE 4-based Kubuntu system.

After inserting a DVD in your drive on a KDE 4-based Kubuntu system, select the K Menu ➪ Applications ➪ Multimedia ➪ Video Player menu item. Dragon will display a dialog with two buttons that enable you to specify whether your want to play a video file or play media such as a DVD). To play a DVD, select the latter, and Dragon will begin playing your DVD, as shown in Figure 19-31.

FIGURE 19-31

Playing a DVD in Dragon



Like most KDE applications, Dragon is quite easy to use. Pressing the F key on your keyboard enables you to switch to full-screen mode and back again. Pressing the R key enables you jump between playing your DVD and its root menu.

Ripping Audio CDs

Ripping is the term that is used to describe the process of extracting the contents of audio CDs and converting them into audio files that you can play in other applications. It isn't necessary to rip a video DVD, because it just contains files in a certain order in certain directories, which you can copy from the command line or from any file manager. Audio CDs are different—they are constructed in a specific format, and contain huge audio files in raw format.

The traditional UNIX philosophy has always been "write small applications that do one thing and do it well." That's a great philosophy for the command line, where you can string a series of applications together to do some amazing thing. That philosophy isn't as applicable or relevant for graphical applications, which are somewhat self-contained by nature. Instead, graphical applications are good examples of the philosophy of convergence, where single applications continue to gain related capabilities. As they grow, these convergent applications can overshadow other graphical applications by making it easier for users to perform a variety of related tasks within a single familiar application, set of menus, and so on. Convergence also makes life easier for the maintainers of a Linux distribution, who can focus on installing and maintaining a single application where two or more might have been required for before.

Why all this philosophy? One casualty of convergence in Ubuntu 8.04 and 8.10 has been Sound Juicer, an excellent CD playing and ripping application that was installed by default in the last few Ubuntu releases. However since Rhythmbox plays audio CDs, Internet radio, podcasts, and a variety of audio file formats, why not just use it to rip CDs, too? Alas, Sound Juicer, I knew thee well. The best part is that Sound Juicer isn't dead, it just isn't installed by default on GNOME-based Ubuntu distributions anymore—devotees can still install up-to-date version of Sound Juicer from the Ubuntu repositories.

The next two sections discuss using Rhythmbox on Ubuntu systems and using k3b on both KDE 3– and KDE 4–based Kubuntu systems.

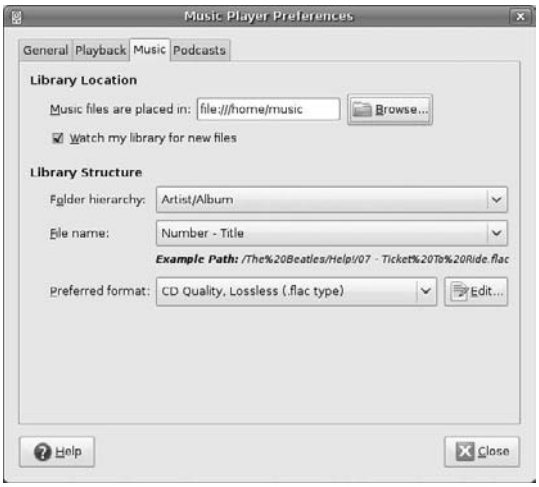
Ripping Audio CDs Using Rhythmbox on Ubuntu

Ripping CDs in Rhythmbox is almost exactly like playing them (as shown in Figure 19-25), with the exception of the fact that you click the Extract button rather than the Play button to start the process. While Rhythmbox is extracting the contents of an audio CD and converting each track to a file, it displays progress information in the lower-left corner of the Rhythmbox dialog.

Before ripping a CD, however, you may want to specify the format of the audio files that Rhythmbox creates, the location in which it creates them, and so on. This is done through Rhythmbox's Preferences dialog. To display this dialog, select the Edit ➪ Preferences menu command. To specify the location to which the audio files that you extract will be saved and the format in which they will be saved, click the Music tab, which is shown in Figure 19-32.

FIGURE 19-32

Rhythmbox preferences for ripping CDs



This dialog enables you to change the following settings:

- **Library Location:** Enables you to select the folder in which you want extracted audio to be stored. Though a folder named Music is automatically created for this purpose in each user's home directory, I tend to create a publicly writable directory called `/home/music` on my systems, to make it easier for multiple users to access a centralized music collection. This also simplifies backups of all of the CDs that all users of your system have extracted. If you want to be really clever, you can either mount a shared network directory as `/home/music` using NFS or share a local `/home/music` directory via Samba so that users of all systems on your local network have access to your music. Ask your users not to copy music from this directory, lest the RIAA sue you for a few million dollars.
- **Library Structure:** Enables you to specify how the audio files that you extract from CDs will be organized in your music library. This section provides the following settings:
 - **Folder Hierarchy:** Enables you to specify the folder hierarchy in which Sound Juicer creates the audio files that it extracts. By default, the tracks of each CD are written in a folder with the name of the CD, which is itself a subdirectory of a folder with the artist's name.
 - **File Name:** Enables you to specify the format for the file name used for each audio track that you extract. The default value, Number - Title, is a good setting to use, because this makes it easy for you to create playlists that follow the original track order of your CDs.
 - **Preferred Format:** Enables you to specify the file format in which extracted audio files are written. Your choices are CD Quality, Lossless (using the FLAC codec), CD Quality, Lossy (using the Ogg container format and Vorbis codec), Voice, Lossless (using the standard WAV file format), and Voice, Lossy (using the Ogg container format and Vorbis codec, but at a lower quality level).

I typically change the location of my music folder and, unfortunately, the Preferred format for any tracks that I extract from my audio CDs. You'll note that the default formats in which you can extract CD audio do not include an option to extract files in MP3 format. This option will only be present if you have installed the `gstreamer0.10-plugins-ugly-multiverse` package (or purchased and installed the commercial Fluendo codecs), as discussed in the "Configuring GStreamer" section of this chapter. If you see an entry for MP3 audio on this menu, thanks for following along earlier. If not, install either the `gstreamer0.10-plugins-ugly-multiverse` package or the commercial Fluendo codecs and restart Rhythmbox to see a menu entry for extracting audio files in MP3 format, which you will probably want to make your default if you want users of other operating systems to be able to play your extracted audio files. The OGG format is far superior to MP3 (and is free of licensing issues), but unfortunately, most people and portable music players can only deal with the MP3 format.

Once you've customized the location of your music files or the format in which they will be extracted, click Close to close the Preferences dialog. You can then click the Extract icon in the Rhythmbox toolbar to begin ripping a CD. While Rhythmbox is extracting the contents of an audio CD and converting each track to a file, it displays progress information in the Status bar at the bottom of the Rhythmbox dialog.

Ripping Audio CDs Using k3b

K3b is installed by default on KDE 3– and KDE 4–based Ubuntu systems, and it is a great application for doing most things with CDs and DVDs. On KDE 3–based systems, starting k3b to rip your CDs is one of the options on the pop-up dialog that appears when you insert an audio CD (Extract Audio with k3b). On KDE 4–based systems, you'll have to start k3b by selecting the K Menu ➪ Applications ➪ Multimedia ➪ CD & DVD Burning menu item. You will then have to select the Tools ➪ Rip Audio CD menu command to put k3b in the right mode to work with audio CDs.

K3b first tries to look up your CD in the freedb.org CD database. If more than one CDDB entry matches your CD or if there is no exact match, you'll see a dialog asking you to select the appropriate entry. K3b will then be populated and display the track listing based on your selection. If your CD can't be found in the database, k3b will use a generic name format such as Track 01, Track 02, and so on.

At this point, your k3b screen will look like the screen shown in Figure 19-33. (The track listing will be different unless you own a fairly obscure CD by the French band "Heldon.")

To proceed, click the Start Ripping button below the title of the CD. K3b displays a dialog that enables you to specify the format in which you want to save the extracted audio files and the directory to which you want to save them, as shown in Figure 19-34.

The audio file formats that are available through this dialog are Wave (raw, lossless PCM audio), Mp3, and Ogg-Vorbis. You can only save audio files in MP3 format if you have previously installed the LAME (LAME Ain't an MP3 Encoder) package—trying to do so without having installed this package will display an error message when you start to rip the CD. Sadly, Mp3 is probably the best format to use if you want to be able to clone your music to your iPod or other portable music player, play your audio files from computers running other operating systems, and so on.

You will probably also want to customize the location to which the extracted audio files are saved, which is typically your `~/Music` directory or some centralized music directory on a file server somewhere that is mounted locally, and so on.

Finally, I also usually select the Create m3u playlist and Use relative paths checkboxes. This will automatically create a playlist file for the extracted audio tracks, which is just a text file that lists the tracks in the order that they appeared on the CD. Checking the Use relative paths checkbox causes the entries in the playlist file to list audio tracks relative to the directory where you stored your extracted audio tracks.

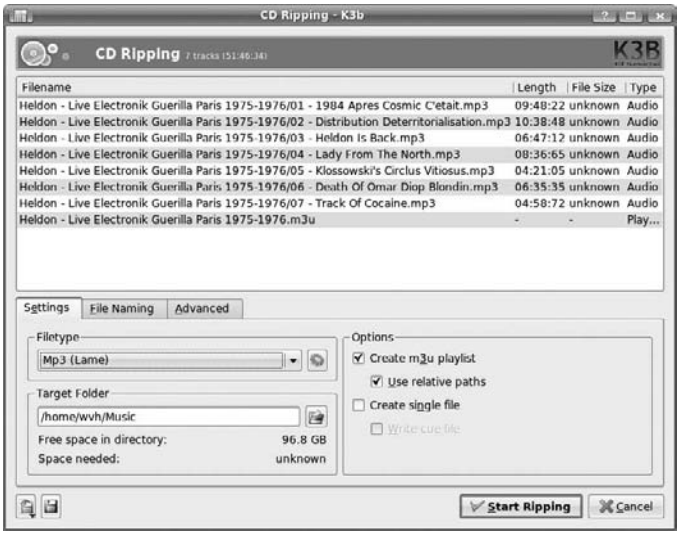
FIGURE 19-33

An audio CD listing in k3b



FIGURE 19-34

k3b preferences for ripping CDs



TIP **Note that the playlist files created by k3b are located in the directory where the directory that holds the extracted audio for a CD is located, not in the same directory with the extracted audio files. This can be confusing if you copy the files to another of your machines, since you may forget to copy the playlist file. Since the default file names for each track of an extracted CD begin with the track numbers, I use a few shell aliases to enable me to quickly create playlist files for a CD. These aren't rocket science, but the ones I use are the following (in my ~/.bashrc file):**

```
alias ftoc="ls *.[Ff][Ll][Aa][Cc] > playlist.m3u"
alias mtoc="ls *.[mM][pP]3 > playlist.m3u"
alias otoc="ls *.[Oo][Gg][Gg] > playlist.m3u"
alias stoc="ls *.[Ss][Hh][Nn] > playlist.m3u"
```

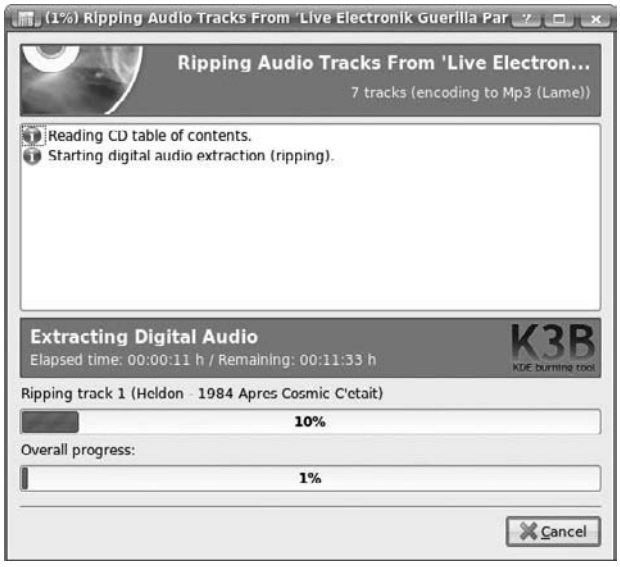
I can then create a playlist file (which I call TOC—Table of Contents—files) called playlist.m3u by changing directory to the directory that holds those files and running one of these aliases. The ftoc alias is for FLAC files, the mtoc alias is for MP3 files, the otoc alias is for Ogg-Vorbis (OGG) files, and the stoc alias is for Shorten (SHN) files.

You can also change the format of the file names that are associated with the extracted audio files by selecting the File Naming tab and modifying its contents, and can use the Advanced tab to change some parameters if you are having trouble reading from an audio CD.

After making any changes that you want to the default k3b settings, click Start Ripping to being extracting audio files form your CD. A status dialog is displayed, as shown in Figure 19-35.

FIGURE 19-35

k3b's status dialog while ripping CDs



A small, general k3b status display dialog is also shown on your screen. This dialog remains on your screen even if you decide to minimize the other k3b dialogs, which is handy to see the status of any k3b task while you're working on something else. If you want to disable this other status dialog, you can do so from the Setting ➪ Configure k3b dialog by selecting the Misc icon and deselecting the Show progress OSD (onscreen display) checkbox, then clicking Apply or OK.

When all of the tracks from your CD have been extracted, a Close button displays on the k3b status dialog shown in Figure 19-35. Click Close to close this dialog and return to k3b. You can exit from k3b by clicking the close box in the window decorations or selecting the File ➪ Quit menu command.

Getting More Information about k3b

I'm not the only fan of k3b—it's a very popular application with many advocates (not all of whom are KDE fanatics—case in point: me). Though I'll be using k3b in other sections of this chapter, it's such a great application that I wanted to provide some general pointers for additional information in case you're interested. Some of the better sites I've found for k3b information, tutorials, and tips and tricks are the following:

- `www.k3b.org`: The k3b application's home page, this site provides a FAQ and access to the latest and greatest releases, as well as the k3b subversion repository if you just can't wait for the next official release.
- `www.novell.com/coolsolutions/feature/2746.html`: A nice article from Novell entitled "CD Burning in Linux Made Easy with k3b," written by Jason Jones.
- `http://sourceforge.net/projects/k3b`: The k3b project's home page at SourceForge.
- `www.newsforge.com/software/03/10/20/1340254.shtml?tid=82`: A nice article and review of k3b from the NewsForge folks, written by Joe Barr.
- `www.pclinuxonline.com/wiki/K3b`: A detailed but informal article from the "PCLinuxOS New User Guide" on using k3b.
- `http://en.wikipedia.org/wiki/K3b`: The k3b entry in the Wikipedia. Cool!

Backing Up DVDs from the Command Line

Ripping CDs is almost a necessity because most people want to be able to play their music files on portable music devices that don't have CD drives. Many fewer people want to extract the contents of DVDs, though anyone in their right mind would probably want to make a backup copy of them. As a long-time UNIX/Linux geek and dog lover, my primary desire in copying DVDs is to back them up so that I still have a copy after our Rottweiler finishes using the original and its case as a chew toy.

Though a variety of different applications are available in the Ubuntu repositories and on the Internet to copy DVDs, most of them have a GUI and take a fair number of mouse clicks to simply copy the DVD. Therefore, my favorite command to back up DVDs is the following:

```
$ dd if=/dev/dvd of=name_of_output_file.iso
```

TIP **You can clone a CD if you want to by executing a similar command, replacing `/dev/dvd` with `/dev/cdrom`, or whatever the appropriate device name is on your system.**

This command is simple, works on every Linux and UNIX system in the known universe, and produces an output file that is an exact image of the original DVD. You can mount the output file just like an actual DVD if you're interested, using commands like the following:

```
$ sudo mkdir /mnt/dvd
$ sudo mount -o loop -t iso9660 name_of_output_file.iso /mnt/dvd
```

The first of these commands creates a directory on which to mount the ISO image; the second actually mounts it and makes it available like any other mounted filesystem.

Once you create an ISO image of a DVD with the dd command, you can also immediately burn the image to a blank DVD-R, as explained later in this chapter.

Creating CDs and DVDs

This section concludes by discussing how to create audio and data CDs, copy CDs and DVDs, burn CDs and DVDs from image files, and assemble files on data or video DVDs. Although standard Ubuntu systems provide a very nice application called Brasero and also support data CD/DVD creation from within the Nautilus file manager, I still personally do all of these things through k3b on both GNOME and Kubuntu systems (though I certainly use Brasero and Nautilus on my laptops, where I want to minimize the amount of software that I install). Nothing personal to the authors of Brasero or Nautilus! Your mileage may vary, so the following sections discuss all of these applications.

Burning CDs Using Brasero on Ubuntu

GNOME's Brasero application makes it easy to create audio CDs from selected audio files on your system or to create data CDs that contain copies of files and directories from your system for backup or portability purposes. This section focuses on creating audio—the steps for creating data CDs are much the same, though you must specify that you are creating a data CD using the Project - New Project ➪ New Data Project menu item.

Creating Audio CDs with Brasero

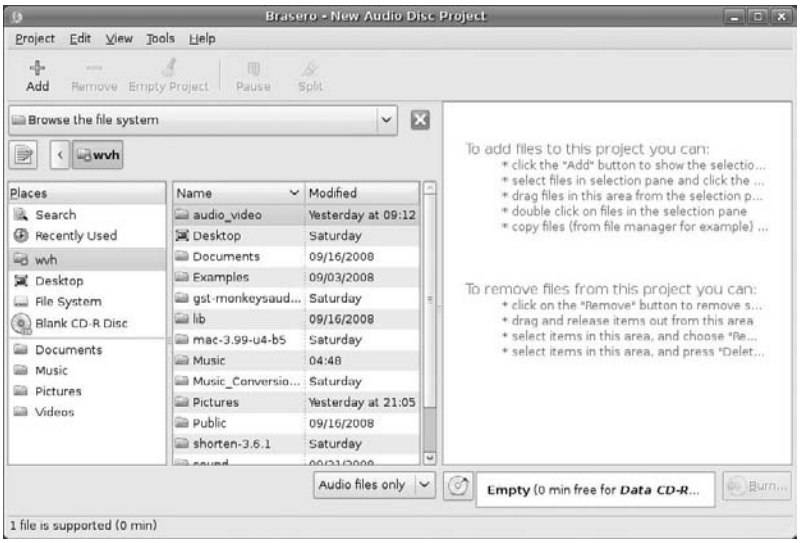
Brasero is one of the application choices that Ubuntu provides in the dialog shown in Figure 19-7 when you insert a blank CD. If you don't see this dialog, you can start Brasero by right-clicking on the Blank CD-R Disk icon and selecting "Open with Brasero Disc Burning," or by selecting the Applications ➪ Sound & Video ➪ Brasero Disc Burning menu item.

Unfortunately, Brasero always starts in Data CD mode, so you'll first need to select the Project ➪ New Project ➪ New Audio Project menu command in order to create an audio CD. This displays the dialog shown in Figure 19-36.

To add files to your new audio CD project, use the browse controls in the left portion of the Brasero screen to navigate to the location where your audio files are stored. Once you display a directory containing audio files that you want to add to your audio CD, you can easily add them to the CD by selecting one or more files and clicking Add. As with most other graphical applications, you can select multiple tracks at one time using Ctrl+left-click or Ctrl+Shift to select a range of songs.

FIGURE 19-36

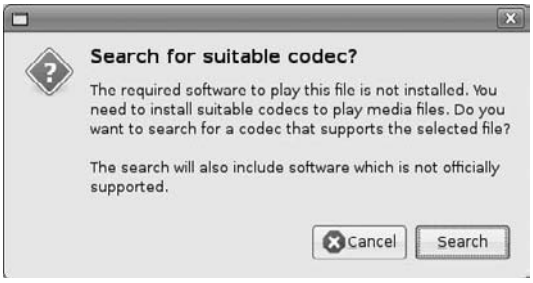
Track selection when burning an audio CD



As you add tracks to your CD, Brasero checks each track to ensure that it knows how to convert a track in that format to the raw audio format that is used on CDs. If it doesn't recognize the format of a track or doesn't have a codec to handle that audio format, Brasero displays an error dialog that says that GStreamer doesn't know how to handle the specified file format, and displays another dialog that offers to look for a suitable codec for you, as shown in Figure 19-37. It doesn't get much more user-friendly than that!

FIGURE 19-37

Ubuntu's initial codec finder dialog



TIP **Seeing the "Find CODEC" dialog on your Ubuntu system generally means that you didn't install the `ubuntu-restricted-extras` package and all of the GStreamer plug-ins that I recommended earlier in this chapter. You may want to do that before proceeding, especially if you're trying to burn CDs from audio files in a variety of formats.**

Click Search to proceed for the search for a relevant codec. If the system can find a suitable codec, it displays a dialog like the one shown in Figure 19-38.

FIGURE 19-38

Selecting a suitable codec for installation



Note that the Install button can't currently be selected—in order to install the specified codec(s), you have to check the checkbox to the left of their names. Selecting this checkbox pops up another dialog, shown in Figure 19-39, that makes sure that you understand that greedy corporate morons have restricted your ability to convert audio files from one format to another in certain countries, and explains the circumstances that would enable you to do so.

FIGURE 19-39

Verifying that it is legal for you to work with your data



After figuring out how one of these circumstances applies to you. click confirm to indicate your grudging acceptance. The dialog shown in Figure 19-38 reappears, this time with the Install button enabled. Click Install to install the missing codec. Installing the missing codec is just like any other software installation process on your system—you'll have to enter your password in order for the package to be installed.

After installing any missing codecs and adding all the tracks that you want to put on your audio CD, the dialog shown in Figure 19-36 reappears, showing the tracks that you've added so far in the right pane of the Brasero window and displaying information about the remaining amount of time available on the CD below that window.

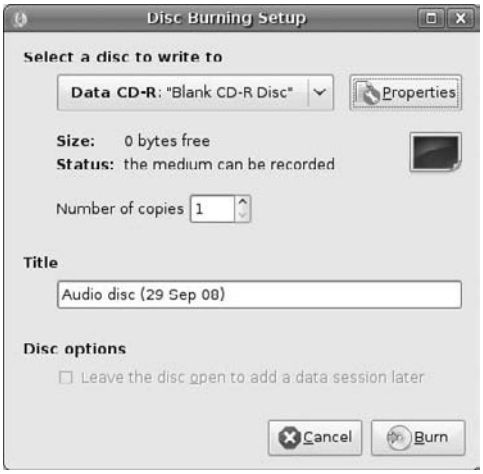
NOTE **Many CD players can only play CDs of a certain size. To maximize compatibility with CD players, click the CD icon below the CD Track Listing window and select the 1 Hour and 20 minutes CD capacity. The 80-minute audio CD is a de facto standard. If you are creating a data CD, you'll see size listings instead—a 650Mb CD is a fine choice for compatibility.**

After adding contents to your CD or DVD, you may want to save the list of the contents of your CD or DVD as a Brasero project, which can be handy if you want to make another, identical CD or DVD at some point in the future. If you don't plan to make another copy, don't bother—but if you might, select the File ➪ Save menu command to display a dialog that enables you to specify the name and location of the saved project file. Brasero's project files are text-format XML files that list the contents of a CD or DVD—they do not actually contain the CD or DVD data.

When you're ready to burn your CD, select the Project ➪ Burn menu command or click Burn in the lower-right corner of the Brasero dialog. The dialog shown in Figure 19-40 appears, requesting confirmation.

FIGURE 19-40

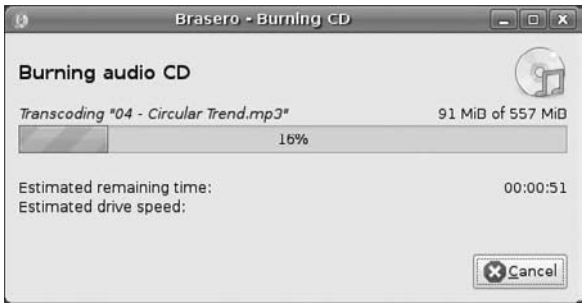
Confirmation request before burning a CD



Click Burn to proceed, and the CD burning operation proceeds. As the process proceeds, Brasero displays a status dialog, as shown in Figure 19-41.

FIGURE 19-41

The progress dialog when burning a CD

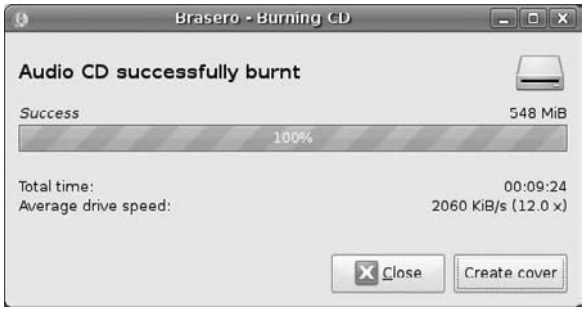


If you are burning an audio CD, Brasero first normalizes the tracks, which means that they will all be at approximately the same volume/equalization (EQ) level. This can be very useful if you are burning a CD that is composed of tracks from multiple sources, each of which probably has its own EQ level.

When writing your CD completes, Brasero ejects the disc and displays a dialog indicating whether writing the disc was successful or not, and giving you the option to create a cover for the CD, as shown in Figure 19-42.

FIGURE 19-42

The success dialog after burning an audio CD



Click Close to close this dialog and return to Brasero. If you want to create a cover for your CD, click Create cover to display Brasero's cover editor, which lets you make simple text covers for your CDs, as shown in Figure 19-43.

If you're finished creating CDs for the moment, you can select the Project ➪ Quit menu command to close Brasero. If you haven't saved the set of files that you were working with as a Brasero project, you'll see a dialog asking if you want to do this before exiting from Brasero.

FIGURE 19-43

Creating a cover for an audio CD



Burning CD and DVD Images Using Brasero

It's quite common to download the boot/live/install CDs for most Linux distributions as ISO (International Standards Organization) images, which means that these are CD or DVD images in a standard format. Brasero makes it easy to create physical CDs or DVDs from these image files. To do so, start Brasero and select the Project ➪ New project ➪ Burn Image menu command. The dialog shown in Figure 19-44 appears.

FIGURE 19-44

Burning a CD or DVD image to disc



Use the Image browse button at the bottom of this dialog to locate the image that you want tot burn to disc, and click Burn. Brasero displays a status dialog much like the one shown in Figure 19-41. When the entire image is written to disc, Brasero displays a confirmation dialog. Click Close to close this dialog and return to Brasero. If you're done creating CDs or DVDs for the moment, you can select the Project ➪ Quit menu command to close Brasero.

Burning Data CDs and DVDs Using Nautilus

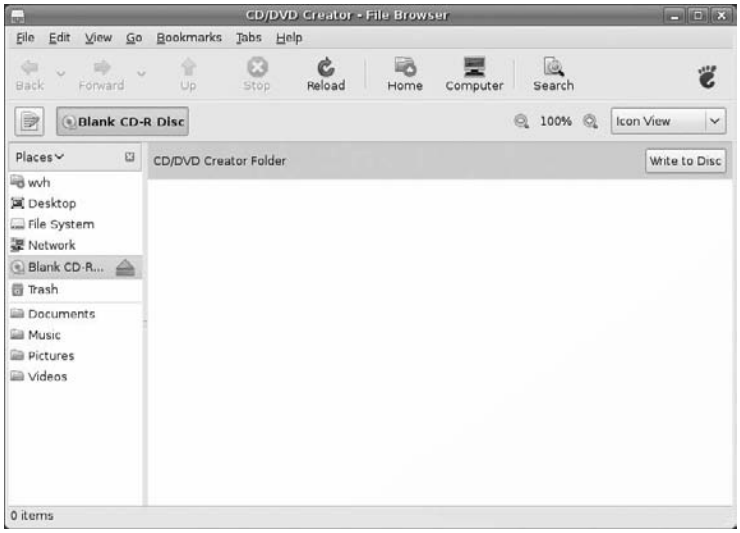
GNOME's CD/DVD Creator application, which is just the Nautilus file manager running in a special mode, is similarly easy to use to create data CDs and DVDs from selected files on your system.

NOTE **You can—only create data CDs and DVDs using Nautilus's CD/DVD Creator—you cannot create audio CDs.**

The CD/DVD Creator is one of the application choices that Ubuntu provides in the dialog shown in Figure 19-7 when you insert a blank CD. (If you don't see this dialog, you can right-click on the Bank CD icon to see your choices.) The CD/DVD Creator's initial screen is shown in Figure 19-45, showing a special virtual folder that is used to track the files and directories that you want to add to your CD or DVD.

FIGURE 19-45

The initial CD/DVD Creator screen

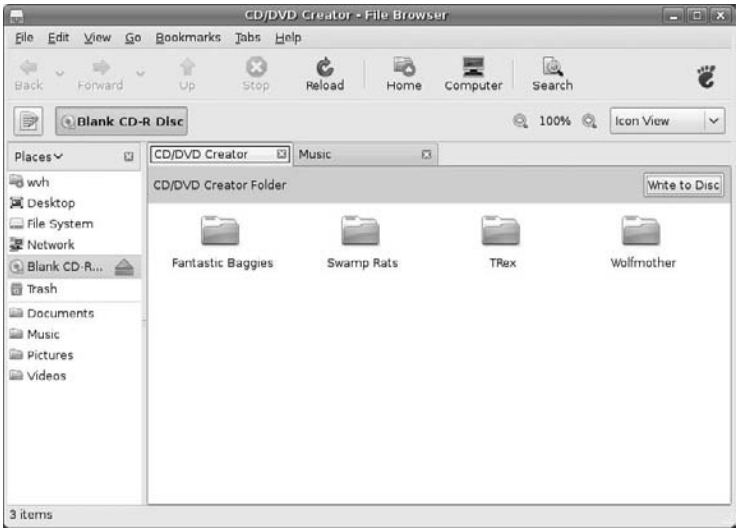


Once this screen displays, you can add files and directories to the CD or DVD by dragging and dropping them from any other location in Nautilus to the Blank CD-R Disc or Blank DVD Disc locations shown in the Places menu at the left side of the Nautilus display. If you're using Ubuntu Hardy (8.04), you can do this by using the single Nautilus window to visit different directories and add file and directories that they contain. If you're using Ubuntu Intrepid (8.10), the new support for tabs in Nautilus makes this even easier because you can visit multiple directories in different tabs and incrementally add files and directories from them (you can see the tabs in Figure 19-40).

When you're done adding files to your CD or DVD, select the Blank CD-R Disc or Blank DVD Disc location shown in the Places menu again to see the files and directories that you've added so far. An updated CD/DVD Creator Folder dialog is shown in Figure 19-46.

FIGURE 19-46

The CD/DVD Creator folder after adding directories



When you're ready to burn your CD, click Write to Disc. The dialog shown in Figure 19-47 appears, requesting confirmation and enabling you to change default values such as the write speed, the name of the CD or DVD, and so on.

FIGURE 19-47

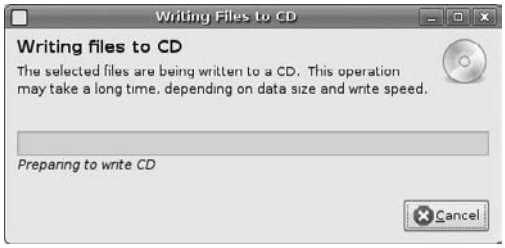
Confirmation request before burning a CD or DVD



Click Write to Disc to proceed, and the burning operation proceeds. As the process proceeds, Nautilus displays a status dialog, as shown in Figure 19-48.

FIGURE 19-48

The progress dialog when burning a CD



When writing your CD or DVD completes, Nautilus ejects the disc and updates this dialog, indicating whether writing the disc was successful or not, and giving you the option to write another disc or close the dialog, as shown in Figure 19-49.

FIGURE 19-49

The progress dialog when burning completes



If you want to burn another CD or DVD with the same contents, simply insert a new blank disc and click Write to Disc. Click Close to close this dialog and return to Nautilus. If you want to burn another audio CD with different contents, you can click Clear to remove the current track listing and start another one. If you're finished creating CDs for the moment, you can select the File ➪ Quit menu command to close Nautilus.

Burning CDs and DVDs Using k3b

K3b is the KDE tool of choice for burning CDs and DVDs on Kubuntu systems (as well as extracting audio from them, as discussed earlier in this chapter). If you're using a KDE 3–based Kubuntu system, you can start k3b from the pop-up dialog that displays when you insert a blank CD or DVD, or by selecting the K Menu ➪ Multimedia ➪ k3b menu command. On KDE 4-based Kubuntu systems, select the K Menu ➪ Applications ➪ Multimedia ➪ CD & DVD Burning menu command. This displays the initial k3b window, as shown in Figure 19-50.

FIGURE 19-50

The initial k3b window

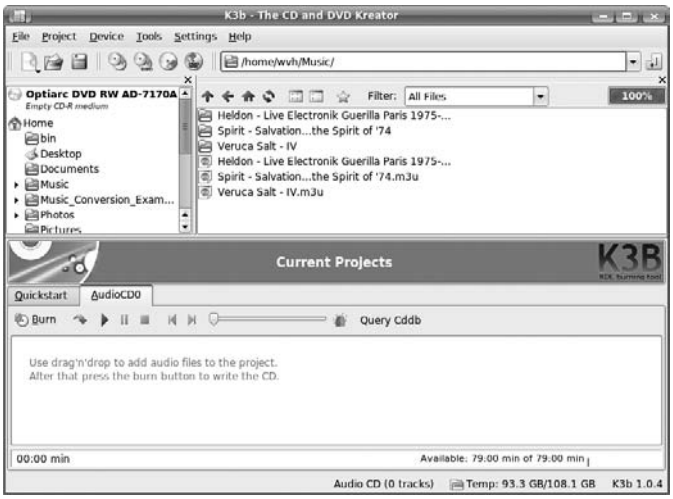


Burning an Audio CD Using k3b

To burn an audio CD from k3b, click the New Audio CD Project button from the main k3b dialog, or select the File ➪ New Project ➪ New Audio CD Project menu command. The dialog shown in Figure 19-51 appears.

FIGURE 19-51

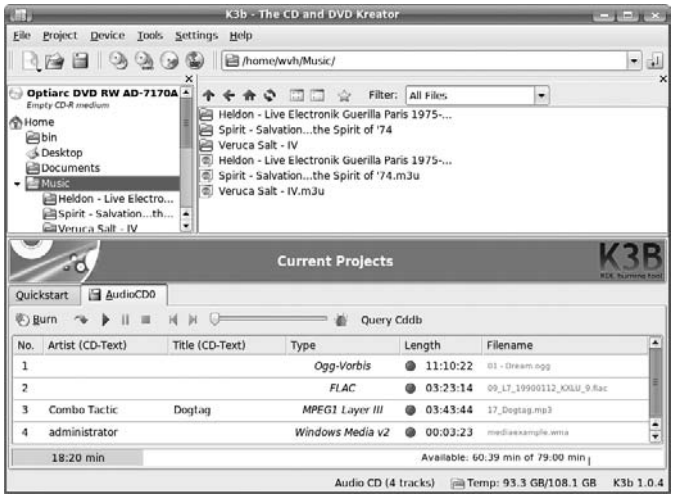
The k3b dialog for an audio CD project



Use the navigation pane at the top right of the screen to navigate to the location(s) where the files that you want to add to your audio CD are located. Selecting a file will add it to the list of files that is slated to be burned to the CD, shown at the bottom of the k3b audio project dialog, as shown in Figure 19-52. As you add audio files, the amount of time used and remaining on the CD displays below the CD contents pane.

FIGURE 19-52

Adding files to an audio CD project



K3b is installed with a number of plug-ins that enable you to add audio files in a number of different audio formats to an audio CD—they will automatically be converted to the raw audio format used on audio CDs. These plug-ins support MP3, MPC, FLAC, OGG, WAV, WMA, and many other audio file formats thanks to plug-ins that leverage the codecs provided by the Sox (`http://sox.sourceforge.net`) and FFMPEG (`http://ffmpeg.mplayerhq.hu`) projects. If you attempt to add a file in a format this is not supported by a k3b plug-in, an error dialog appears that informs you that the file could not be added to the project.

Once you have added all of the audio files that you want to add to your CD, click Burn to display a dialog that enables you to customize many aspects of the CD that you are creating, as shown in Figure 19-53.

The tabs of this dialog enable you to customize a disk image, add CD-TEXT to your CD, and do some other advanced customization. One option that I typically select is the Advanced tab's Normalize Volume Levels checkbox, which records all of the tracks on the CD at approximately the same volume equalization (EQ) level. (Checking this option requires that you install the `normalize-audio` package on your Kubuntu system—it isn't installed by default.)

After making any customizations that you want, click Burn to begin writing your audio CD. As your CD is created, k3b displays the status dialog shown in Figure 19-54.

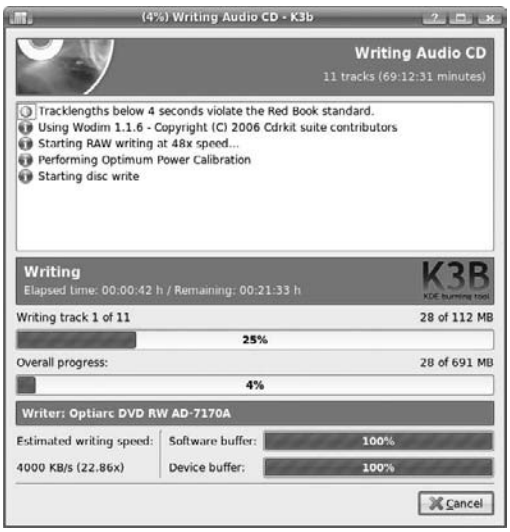
FIGURE 19-53

Getting ready to burn an audio CD



FIGURE 19-54

Burning an audio CD



Once writing completes, k3b displays a final status message. To return to the main k3b dialog, click Close. Eject your CD and have fun listening to it!

Burning an ISO Image to CD or DVD with k3b

To burn an ISO image to a CD or DVD after starting k3b, click the Burn CD Image or Burn DVD ISO Image buttons from the main k3b dialog, or select the Tools ➪ Burn CD Image or Tools ➪ Burn DVD ISO Image menu commands. I'll use burning a DVD image as an example in this section, and therefore the dialog shown in Figure 19-55 is displayed.

FIGURE 19-55

The k3b dialog for burning an ISO image



To navigate to the directory where your ISO image is located, click the folder icon at the right of the Image to Burn text box. A standard KDE navigation dialog appears, as shown in Figure 19-56. Once you've selected your ISO image, click OK to close the navigation dialog. k3b checks the integrity of the ISO image file and calculates an MD5 checksum to use when verifying the copy.

To begin writing the image to a DVD, insert a blank, writable DVD into the specified drive and click Start. A dialog like the one shown in Figure 19-57 appears as k3b fills its internal buffers and begins the write process.

Once writing completes, k3b displays a status message. To return to the main k3b dialog, click Close.

FIGURE 19-56

Selecting a DVD ISO image in k3b

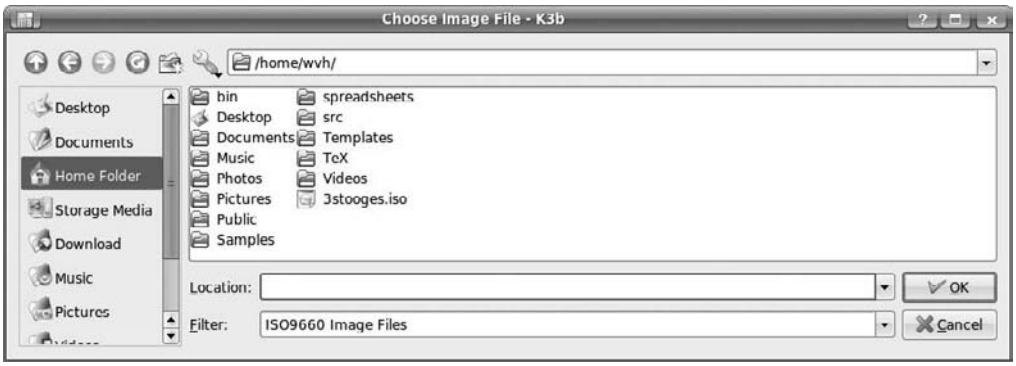
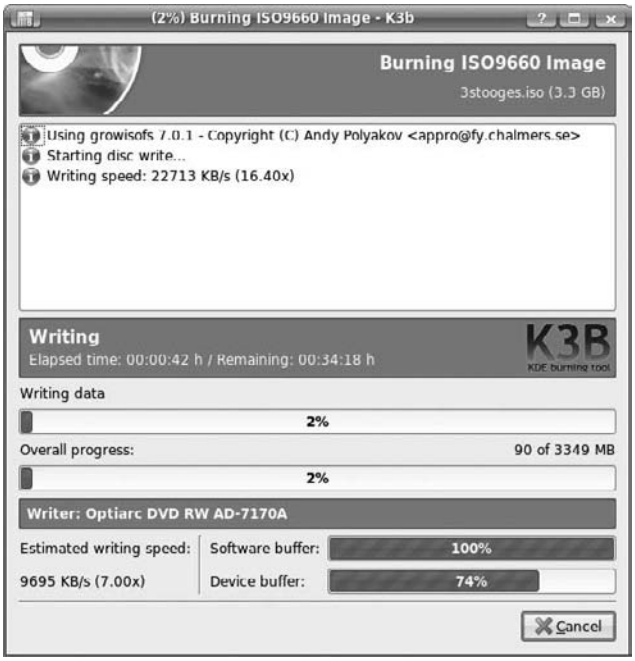


FIGURE 19-57

Status information while burning an ISO image to DVD in k3b



Copying a CD or DVD with k3b

The k3b CD or DVD copy process consists of several steps: k3b first creates an ISO image of the CD or DVD, ejects the original medium, requests a new CD or DVD, and then writes the ISO image to the new CD or DVD.

To copy a CD or DVD after starting k3b, select either the Tools ➪ Copy CD or Tools ➪ Copy DVD menu commands. I'll use burning a DVD image as an example in this section, and therefore the dialog shown in Figure 19-58 is displayed.

After verifying that the source and destination devices for reading and writing the DVD (respectively) are correct, click Start to begin the copy process. The dialog shown in Figure 19-59 is displayed.

FIGURE 19-58

The DVD Copy dialog in k3b



After k3b reads the disc and creates its ISO file, it will display a request for DVD insertion if the source and destination drives are the same, as shown in Figure 19-60.

Insert the new disc and wait a few seconds. The media request dialog closes, and k3b writes its ISO file to the blank DVD. Voilà, an exact copy!

NOTE **KDE 3–based Kubuntu systems may pop up the standard media insertion dialog, asking what you want to do with the new DVD. Just click Cancel to close this dialog.**

Once writing completes, k3b displays a status message to that effect in the Copy DVD dialog. To return to the main k3b dialog, click Close.

FIGURE 19-59

Status information while copying a DVD in k3b



FIGURE 19-60

A request for a blank disc when copying a DVD in k3b



Working with Other Audio Sources

Thanks to the Internet, portable music players, podcasting, and greater amounts of local disk space to devote to online audio files, playing audio on your computer isn't just about playing CDs anymore. A huge variety of alternate audio sources are available today, including streaming broadcasts over the Internet, downloaded podcasts (which are nothing more than MP3 files), and ripped audio files (so that you can store your original CDs safely away somewhere as backups). This section explains how to use Ubuntu's Rhythmbox and Kubuntu Amarok applications to play audio in sources that it understands, and how to use Ubuntu's excellent Sound Converter application to convert audio files from their original formats into formats that you can more easily use on your computer system.

Playing Audio Files and Internet Radio Using Rhythmbox

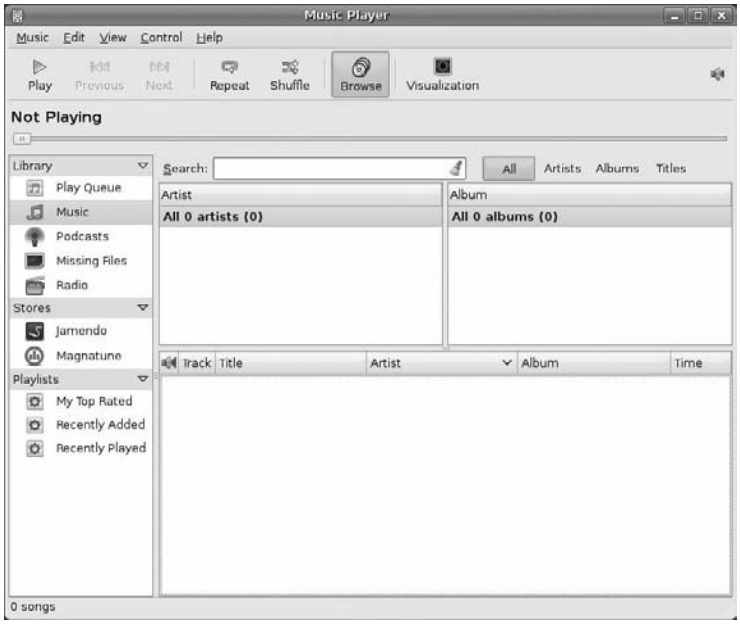
Rhythmbox was discussed earlier in this chapter in the context of playing CDs, but it's really much more. Rhythmbox is oriented toward audio delivered in a variety of file formats, whether as a CD, a stream that you can read over the Internet, or a physical file that you downloaded or ripped, stored on your computer system, and subsequently want to play.

To start Rhythmbox, select the Applications ➪ Sound and Video ➪ Rhythmbox Music Player menu command.

Figure 19-61 shows the default Rhythmbox interface with no audio libraries, podcasts, or other audio sources specified.

FIGURE 19-61

Rhythmbox with no imported audio sources



Next, you may want to add information about the location of your audio files if you haven't already done so when ripping CDs. To do this, select the Edit ➪ Preferences menu command. Click the Music tab of the dialog, which is shown in Figure 19-62.

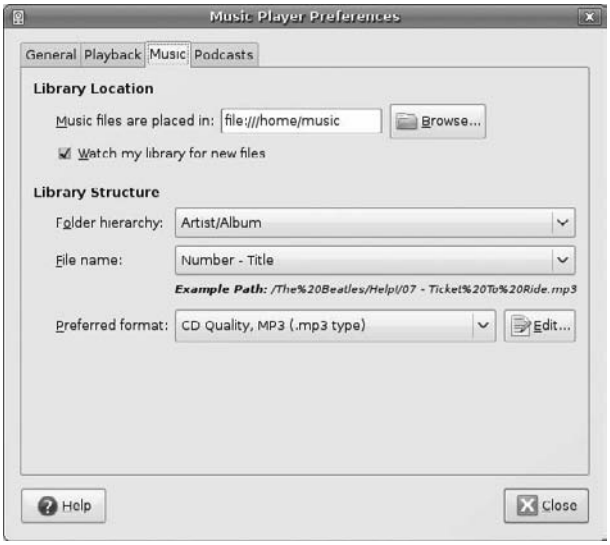
To change the name of the directory in which Rhythmbox will look for your audio files, click the Library Location's Browse button and navigate to that folder in the resulting dialog.

Once you've specified the directory where your music files are located, select the Watch my library for new files checkbox to ensure that Rhythmbox notices any new audio files that you create or put in that directory, and click Close to exit the Preferences dialog. Rhythmbox begins the import process, updating its display as new files are added to its list of available music files. Any problems that Rhythmbox

encounters during this process are displayed in a separate window so that you can examine them. In most cases, these are file corruption problems, but at least you can delete the offending files in your library and re-extract them from your original CDs. Figure 19-63 shows Rhythmbox after importing the contents of a subset of my music library.

FIGURE 19-62

Identifying the location of music files in Rhythmbox



Once you've told Rhythmbox about your library, you can see that it creates summary entries for each artist across all of the directories in your library, and also creates per-directory listings that correspond to a playlist for each Album/CD in your library. To play a particular Album, select its name from the Album list and click Play. To see all of the Albums available from a specific artist, click the name of that artist in the Artist window. You can then select the Album that you want to listen to and click Play. Once you're playing a specific Album, you can also skip from track to track by clicking the name of the track that you want to play in the track listing

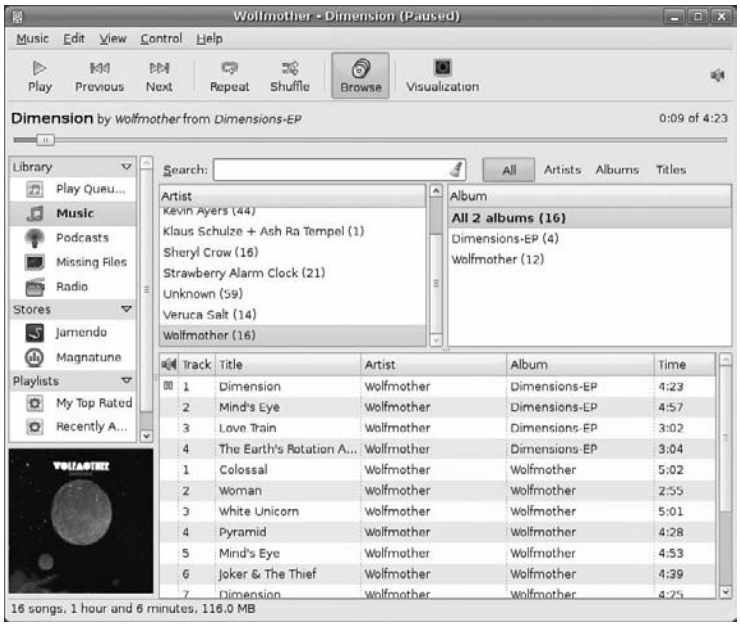
Creating Playlists

If you're like me, you may own a lot of music, but only truly love a few tracks on each album. Rhythmbox makes it easy for you to create your own "Greatest Hits" collections by making it easy to create playlists that contain specific songs from any of the Albums in your library.

To create a new playlist and add songs to it, select the Music ➪ Playlist ➪ New Playlist menu command (or type Ctrl+N). A new playlist is displayed at the bottom of the Source pane at the left of the Rhythmbox window. To assign it a specific name, left-click its blank name and enter the name that you want the new playlist to have. To change the name later, right-click on its name and select Rename from the pop-up menu. You can then type its name and press Return to give it the new name.

FIGURE 19-63

Rhythmbox after loading music files from a library



To add songs to a playlist, simply browse your library and drag and drop the tracks that you want to add to the name of the playlist that you want to add them to in the sources window. Figure 19-64 shows the contents of a playlist containing my favorite tracks by the artist Kevin Ayers. My favorite tracks today that is—one of the great things about creating your own playlists is that you can easily add and remove tracks from them to match your current mood or favorite tracks.

To remove a song from a playlist, right-click on its name and select remove from the pop-up menu. This only removes the song from the playlist, not from your library—you can always add it to the playlist again if you want to hear it again.

Playing and Adding Internet Radio Stations

Internet radio is a revolutionary step for broadcast media because it enables people anywhere to listen to their favorite radio stations regardless of whether you can physically receive their signal (or even have a radio, for that matter). Whether or not the RIAA will succeed in killing off music via Internet radio is still up in the air—but let's enjoy it while we can!

To listen to an Internet Radio station, you can simply click the Radio entry in the Rhythmbox Source pane and double-click the name of the station that you want to listen to. To listen to a station that isn't present in Rhythmbox's default list of stations, right-click the Radio entry in the Rhythmbox Source pane, and select New Internet Radio Station from the pop-up menu. A dialog appears in which you can enter the URL of the Internet radio station that you're interested in. Click Add to add that station to your list of current radio stations, as shown in Figure 19-65.

FIGURE 19-64

A custom playlist in Rhythmbox

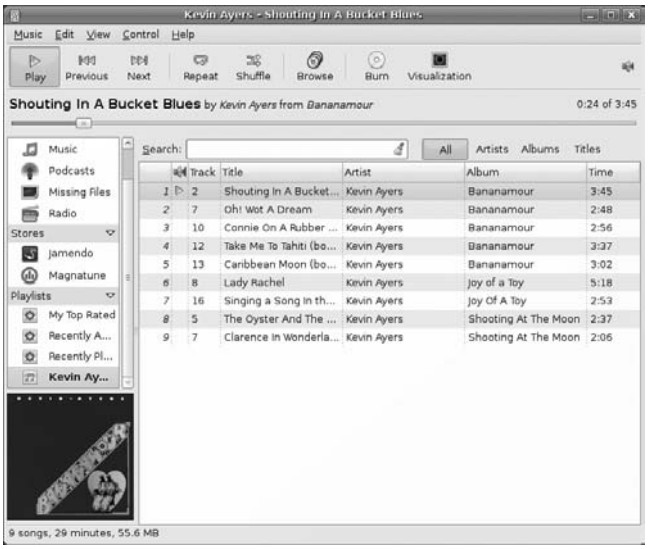
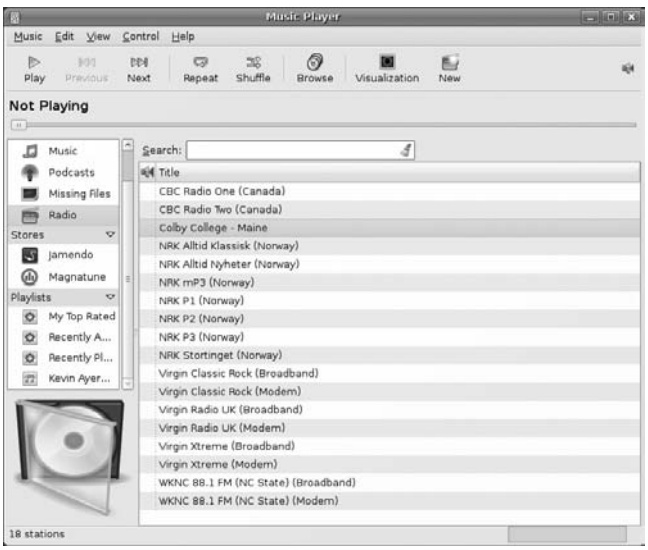


FIGURE 19-65

Playing Internet radio in Rhythmbox



You can then click on the Radio entry to show all of your Internet radio stations, and listen to the station you're interested in by simply double-clicking on its name.

TIP **By default, radio stations identified by URL show up as URLs in the radio station list—you can right-click on an entry and select Properties from the pop-up menu to edit how the radio station appears in the Radio list.**

Adding and Playing Podcasts

Podcasts are revolutionizing broadcast media because they give everyone the opportunity to share their opinions with anyone else. Podcasts make it easy for you to subscribe to certain shows and listen to them at your leisure. Podcasts are especially attractive for people on the move because they provide a great way for people anywhere to keep up with news, church sermons, and other audio presentations from their hometown.

To subscribe to a podcast, right-click the Podcast entry in the Rhythmbox Source pane, and select New Podcast feed from the pop-up menu. A dialog appears in which you can enter the URL of the feed for a podcast you're interested in. Click Add to add that feed to your list of current Podcast feeds, as shown in Figure 19-66.

You can then click on the Podcast entry to show all of your Podcast feeds, and download any available podcast by right-clicking on an available podcast and clicking Download episode from the pop-up-menu. Once a podcast has downloaded, you can listen to it by simply double-clicking on its name.

You can customize where your downloaded podcasts are stored and the frequency with which Rhythmbox looks for new podcasts by selecting the Edit ➪ Preferences menu command and clicking the Podcast tab. This displays the dialog shown in Figure 19-67, which enables you to update those settings.

FIGURE 19-66

Downloading and playing podcasts in Rhythmbox

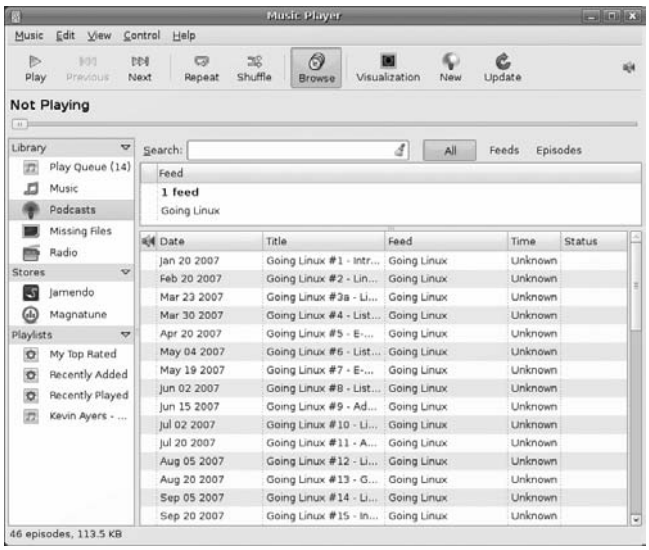
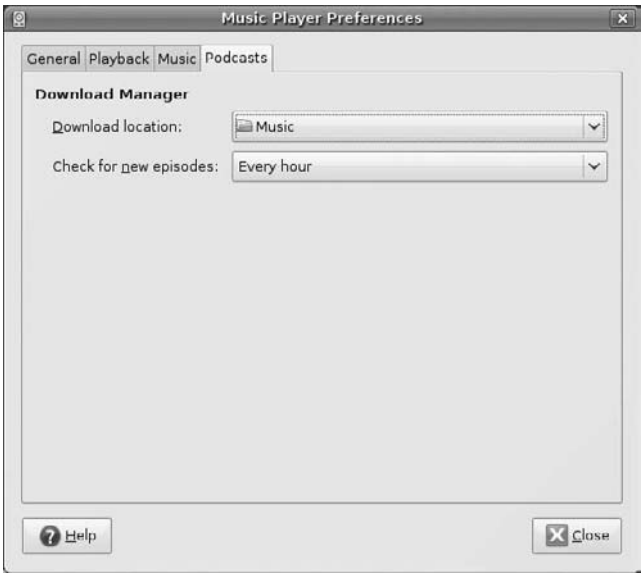


FIGURE 19-67

Customizing podcast settings in Rhythmbox



Playing Audio Files and Internet Radio on Kubuntu

Amarok was discussed earlier in this chapter in the context of playing CDs, but it's really much more. Amarok is oriented toward audio delivered in a variety of file formats, whether as a CD, a stream that you can read over the Internet, or a physical file that you downloaded or ripped, stored on your computer system, and subsequently want to play.

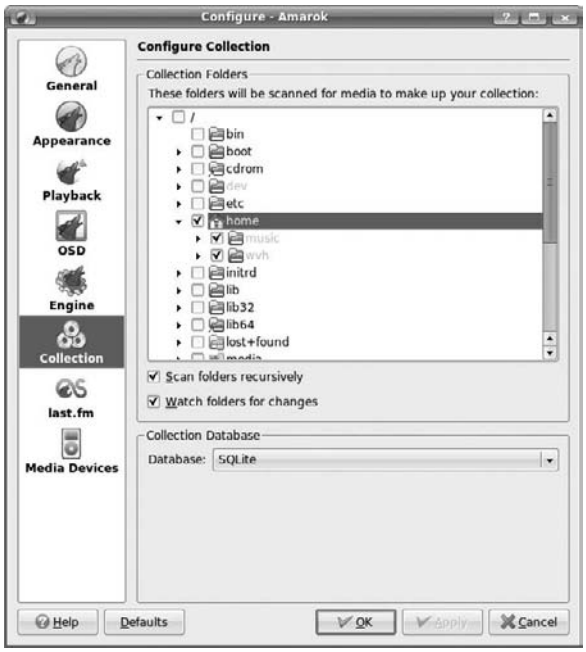
To start Amarok on a KDE 3-based system, select the K Menu ➪ Multimedia ➪ Amarok menu command. To start Amarok on a KDE 4-based system, select the K Menu ➪ Applications ➪ Multimedia ➪ Amarok menu command.

Amarok provides a number of tabs along its left-hand side that enable you to access various features. The primary tabs that are of interest here are the Collection tab, which provides access to your audio files, and the Playlists tabs, which provides access to playlists that enable you to access various audio sources. Playlists are essentially files that list audio resources that you can access and, if multiple audio sources are listed, the order in which those resources should be played. Typically, each CD (or album, for you Luddites) has an associated playlist that defines the track listing on the original medium.

When you first start Amarok, the Context tab appears, prompting you to click a button to build a collection. A collection is just a list of the current audio files that are found on your system. Unlike some audio players, Amarok simply builds its own internal listing of the audio files that it finds—it doesn't import them into any special database or other format. To specify the location(s) that Amarok will search for audio resources, select the Settings ➪ Configure Amarok menu command to display Amarok's Configuration dialog, and click the Collection icon in the left-hand menu. The dialog shown in Figure 19-68 appears.

FIGURE 19-68

Specifying audio file locations in Amarok

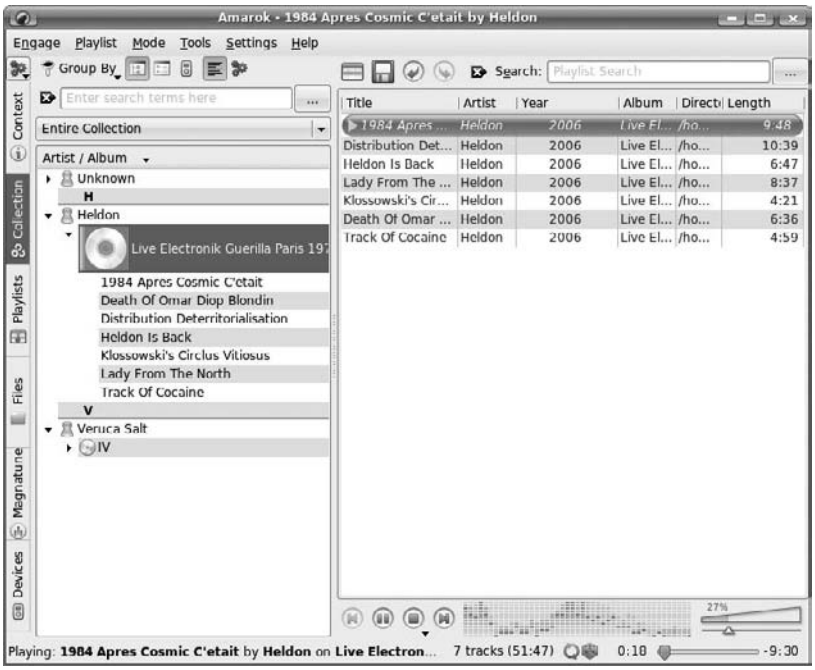


Click the checkbox beside each directory that you want Amarok to search for audio files. You can drill down into directories by clicking the arrow to the left of the checkbox beside their names. You will probably want to select the Music subdirectory of your home directory, as well as any common directories in which you rip CDs or store audio files from other sources. The Scan folders recursively and Watch folders for changes checkboxes, located below the directory listing, are selected by default. The first of these can be useful if, for example, your system has multiple users (and you can access their directories)—leaving this checkbox selected and identifying `/home` as a directory that you want to scan will automatically locate any audio files that any of your other users store in their home directories. The Watch folders for changes checkbox is similarly useful, because it enables Amarok to automatically monitor all selected folders and automatically add any new audio files that it finds to your collection.

Once you're finished specifying audio file locations, click OK to close this dialog. You can then build your initial collection. Once this has finished, click the Collection tab to see the audio files that Amarok has found. Each directory of audio files, which usually represent the audio files associated with a specific artist, will be represented as a top-level entry in the listing pane. To play any of these, click on any entry to drill down into its subdirectories, which generally represent files that have been extracted from specific CDs. To play the files for a given entry, right-click on its name and select Load from the context-sensitive menu that appears. The audio files associated with that entry display in the right-hand pane and should start playing immediately. The right pane is known as the Playlist pane, because it shows the list of the current tracks that will be played, in order. To play a specific track, double-click its name in the listing pane. Figure 19-69 shows Amarok playing the files that I extracted from a specific CD by one of my favorite artists, the French band Heldon.

FIGURE 19-69

Playing a specific recording from a collection



Creating Playlists

Amarok also makes it easy for you to create your own sets of songs or other audio tracks from multiple sources and save them. These listings of audio tracks that you want to play in a given sequence are known as playlists. For example, suppose that you always wanted to hear the audio files from two CDs, one after the other. To create a playlist that does just that, load the playlist from one entry in your collection, then navigate to another entry in your collection, right-click on its name, and select Append to playlist from the context-sensitive menu that displays. The tracks from the newly selected collection entry will be appended to the list that is current shown in the Playlist pane. To save this list as a playlist, click the floppy disk icon above the Playlist pane, enter the name for this playlist, and click OK. Your new playlist is added to the list of your personal playlists, which is saved in the screen that you see when you click Amarok's Playlist tab, and that tab is displayed, as shown in Figure 19-70.

Amarok's Playlist tab provides access to much more than the playlists that you create and save. The Smart Playlists entry in the Listing pane, shown expanded in Figure 19-71, provides a number of automatic playlists that you can use to play tracks from your audio collection in different ways.

One of my favorites of these is the 50 Random Tracks playlist, which randomly selects 50 tracks from your collection and plays them. This is much like the Shuffle mode offered by many portable media players, and gives you the chance to hear selections from your audio collection that you might ordinarily think to select and play manually.

FIGURE 19-70

Viewing playlists in Amarok's Playlist tab

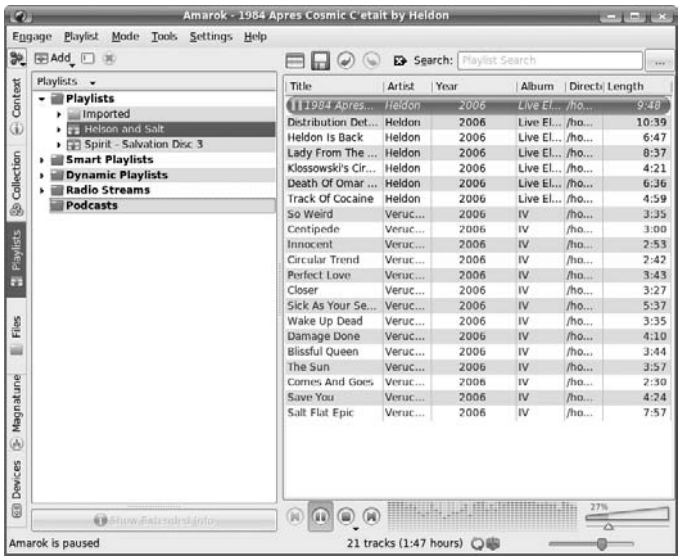
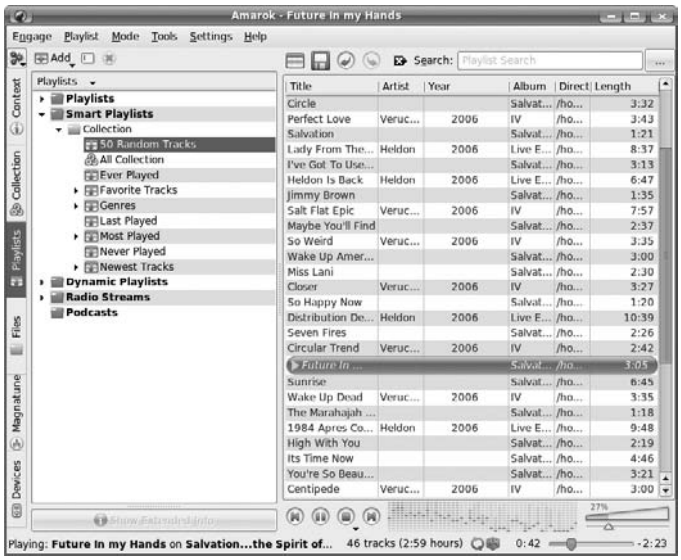


FIGURE 19-71

Amarok's Smart Playlists



Playing and Adding Internet Radio Stations

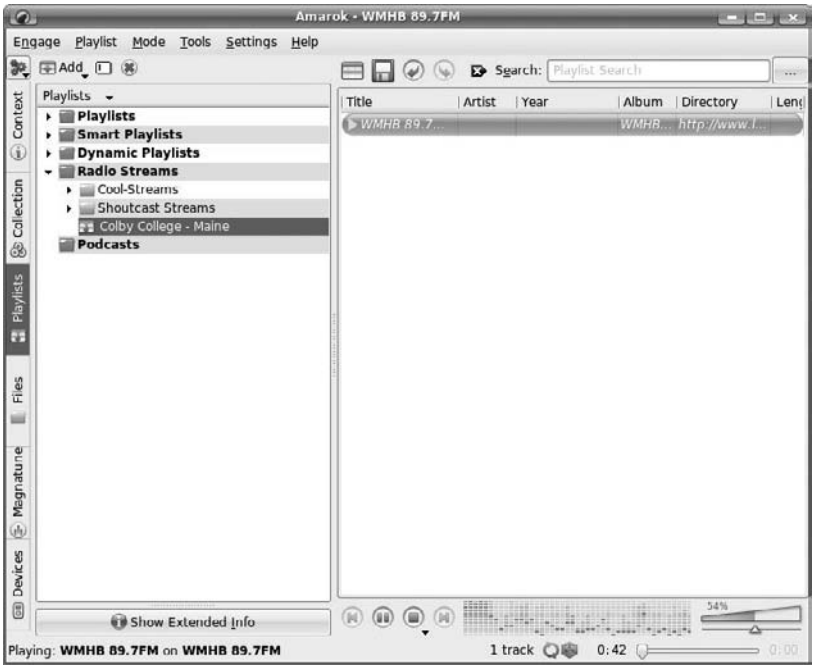
Internet radio is a revolutionary step for broadcast media because it enables people anywhere to listen to their favorite radio stations regardless of whether they can physically receive their signal (or even have a radio, for that matter). Whether or not the RIAA will succeed in killing off music via Internet radio is still up in the air—but let's enjoy it while we can!

Amarok's Playlist tab makes it easy for you to listen to an Internet radio station. To do this, expand the Radio Streams entry to see the list of predefined Internet radio stations that are available from the Internet radio sources Cool-Streams (`www.coolstreams.com`) and Shoutcast (`www.shoutcast.com`). To select and play an entry from one of these, double-click the name of the station that you want to listen to. Amarok retrieves information from the Internet location of that radio station and begins playing the station. It really couldn't get any easier.

To listen to a station that isn't present in Amarok's default list of stations, right-click the Radio Stream entry in the Rhythmbox Listing pane, and select New Radio Stream from the pop-up menu. A dialog appears, in which you can enter the URL of the Internet radio station that you're interested in. Click Add to add that station to your list of current radio stations. To play the newly added station, click the Stop icon below the Playlist pane to stop anything that's currently playing, and double-click on the newly added station. The new station will then begin to play, as shown in Figure 19-72.

FIGURE 19-72

Playing a new Internet radio station in Amarok



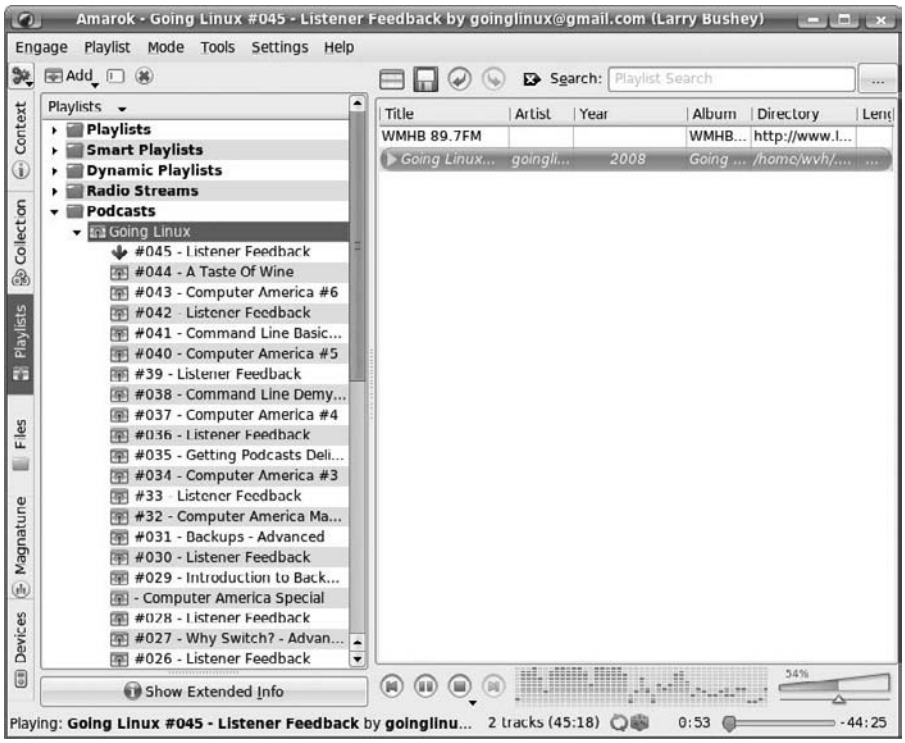
Adding Podcasts

Podcasts are revolutionizing broadcast media because they give everyone the opportunity to share their opinions with anyone else. Podcasts make it easy for you to subscribe to certain shows and listen to them at your leisure. Podcasts are especially attractive for people on the move because they provide a great way for people anywhere to keep up with news, church sermons, and other audio presentations from their hometown.

To subscribe to a Podcast, right-click the Podcast entry in Amarok's Listing pane, and select Add Podcast from the pop-up menu. A dialog appears, in which you can enter the URL of the feed for a podcast you're interested in. Click Add to add that feed to your list of current Podcast feeds, as shown in Figure 19-73.

FIGURE 19-73

Podcast feeds in Amarok



You can then click on the Podcast entry to show all of your Podcast feeds, and download any available podcast by right-clicking on an available podcast and clicking Download Media from the pop-up-menu. Once a podcast has downloaded, you can listen to it by simply double-clicking on its name.

You can check for updates to an existing Podcast entry by right-clicking on its name in the Listing pane and selecting the Check for Updates menu entry.

Converting Audio File Formats

It's almost inevitable that you'll eventually get some great music from free downloads, off the Internet, or after ripping it from your CDs, only to find that it isn't in "the right" format for some other purpose. The best examples of this are things like:

- Ripping CDs into Ogg Vorbis format and then wanting to play them on an iPod, which can't play OGG files.
- Downloading live concert recordings in FLAC format, only to find that this format is not usable with today's iPod (and probably never will be because of the size of the FLAC files).
- Wanting to create a greatest hits CD for your car, only to find that you can't create a playable audio CD from MP3, OGG, or FLAC files.

Regardless of which corner you find yourself in, it's irritating to have music that you can play on your desktop system but that you can't use in portable music players, to burn CDs, and so on. The problem that these examples expose is the fact that specific devices and types of media require specific audio formats. Apple's iPods can only play files in MP3 or Apple's own formats; audio CDs must contain WAV files, and so on. CD-burning applications such as Brasero on Ubuntu systems and k3b on Kubuntu systems will automatically convert files to the WAV format that is required when burning CDs, but they do so on the fly. What if you actually want to convert audio files that you already have so that you can permanently use them in some other format?

Luckily, both Ubuntu and Kubuntu provide a great solution to these sorts of problems through applications called Sound Converter (Ubuntu) and soundKonverter (Kubuntu—notice the "K"). As the names imply, these applications convert digital audio files from one format to another.

Installing and Using Sound Converter on Ubuntu Systems

Because not everyone will be using their Ubuntu system for serious audio work, Sound Converter isn't installed by default. You must install Sound Converter yourself from the `soundconverter` package, using your favorite package management application as discussed in Chapter 23.

TIP **Sound Converter uses the GStreamer framework discussed earlier in this chapter in the section on "Installing the GStreamer Framework and Plug-Ins." You must follow the instructions in that section before trying to use Sound Converter—well, at least before trying to successfully use Sound Converter.**

Once you've installed Sound Converter, you can execute it by selecting the Applications ➪ Sound & Video ➪ Sound Converter menu command. Sound Converter's primary dialog displays, as shown in Figure 19-74.

To add the files that you want to convert, click Add File to display a dialog that enables you to navigate to and select the files which you want to convert to another format. Once you've selected those files, click Open to add the selected files to the list of files that you want to convert and close the navigation dialog. The files that you selected display in the main Sound Converter dialog, as shown in Figure 19-75. (You can also click Add Directory to select and add an entire directory of files.)

Next, specify your format conversion preferences. The files that I've selected in Figure 19-75 are in various file formats, but my goal is to be able to play all of these on my iPod, which requires MP3 files. Not a problem—I just need to specify MP3 as the target output format. To do this, select the Edit ➪ Preferences menu command to display the dialog shown in Figure 19-76.

FIGURE 19-74

Starting Sound Converter



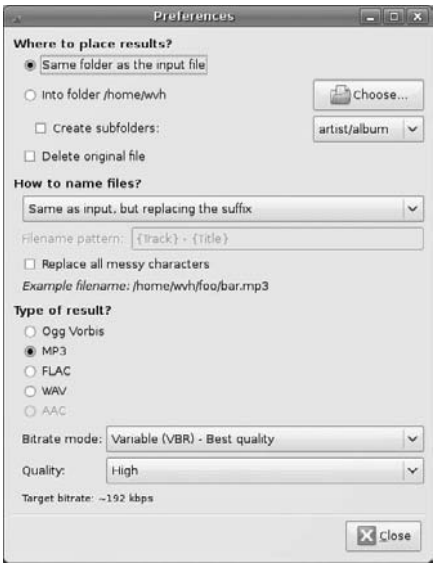
FIGURE 19-75

Files slated for conversion



FIGURE 19-76

Specifying format conversion preferences



The Where to place results? section of Figure 19-76 determines where the converted output files are located. (I tend to leave them in the same directory so that I know what they are.) The Type of Result? section of this dialog shows that I am converting the files I selected to MP3 format, as desired. Click Close to exit the Preferences dialog. Once this dialog closes, you can click the Convert icon to begin the conversion process.

Once the format conversion is complete, the converted files are ready to be copied to your target device, or to do whatever else you want to do with them.

TIP **One of the few audio formats that isn't somehow supported by the GStreamer framework is Shorten (SHN), a free lossless audio codec that is quite popular in live recording circles. To work with files of this type, you can build and install the XMMS audio player and a plug-in for XMMS so that you can play them directly (as described in the example of building open source software in Bonus Chapter 3), or you can use your favorite package management software to install the `shntool` package, a command-line utility that can convert them to WAV files so that Sound Converter can convert to your favorite output format.**

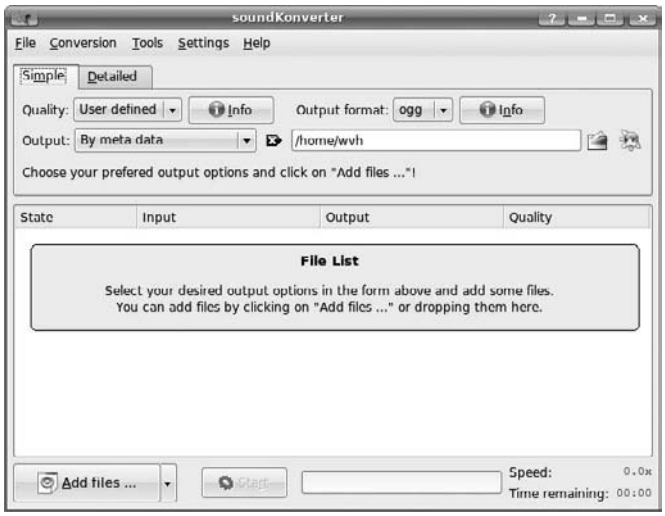
Installing and Using soundKonverter on Kubuntu Systems

Because not everyone will be using their Kubuntu system for serious audio work, soundKonverter isn't installed by default. You must install soundKonverter yourself from the `soundkonverter` package, using your favorite package management application as discussed in Chapter 23.

Once you've installed soundKonverter, you can execute it by selecting the K Menu ➪ Applications ➪ Multimedia ➪ soundKonverter menu command on KDE 3-based systems, and the K Menu ➪ Applications ➪ Multimedia ➪ soundKonverter menu command on KDE 4-based systems. Its primary dialog appears, as shown in Figure 19-77.

FIGURE 19-77

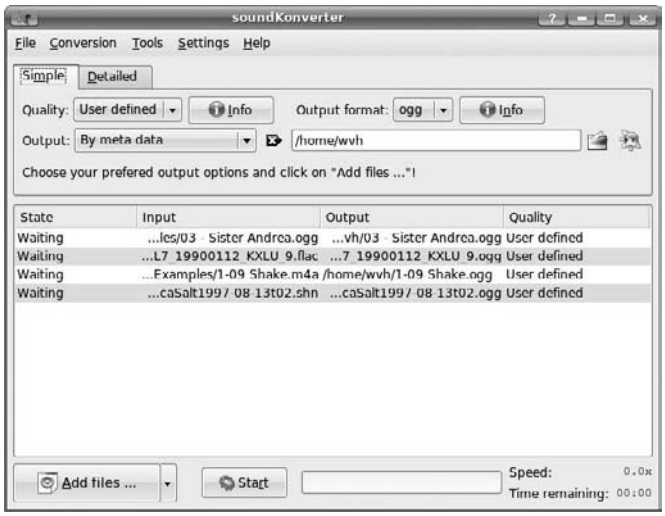
Starting Sound Converter



To add the files that you want to convert, click Add Files to display a dialog that enables you to navigate to and select the files which you want to convert to another format. Once you've selected those files, click Open to add the selected files to the list of files that you want to convert and close the navigation dialog. The files that you selected display in the main soundKonverter dialog, as shown in Figure 19-78.

FIGURE 19-78

Files slated for conversion



Next, specify your format conversion preferences. The files that I've selected in Figure 19-78 are in various file formats, but my goal is to be able to play all of these on my iPod, which requires MP3 files. Not a problem—I just need to specify MP3 as the target output format. To do this, click the drop-down Output format menu to display available formats, and select MP3. Next, you'll probably want to specify how soundKonverter figures out what to call the converted files. This is controlled by the options on the Output drop-down menu, which by default is set to "By meta data." This setting is useful if all of the files that you are converting contain correct title information, which is stored in tags in the headings used of many types of audio files. Unfortunately for me, this isn't always the case in the types of files that I want to convert, so I typically set this to "Copy directory structure," which makes a copy of the directory containing the input files and names the output files based on their original files names, but with a new file extension that reflects their format. I typically create a soundKonverter directory in my home directory and click the folder icon to navigate to that directory and set it as my output directory.

Once you've finished changing any of the default settings, click the Start button to begin the conversion process. Once the format conversion is complete, the converted files are ready to be copied to your target device, or to do whatever else you want to do with them.

Summary

This chapter covered a lot of general audio and digital media ground to help you get the most out of your Ubuntu or Kubuntu system. It began with a Rosetta Stone for the most common terms and acronyms associated with digital audio and video, whether stored on a CD or DVD, or in files on your computer. Subsequent sections explored how to configure the bleeps, boops, and boings your system can emit when various events occur, and explored the GStreamer framework that provides the underpinnings of all GNOME-based audio and video applications, explaining the packages that you'll want to install to deal with the zillions of audio file formats used on personal computer systems today. Other sections discussed the new audio systems used on Ubuntu and Kubuntu systems; how to play, rip, and burn CDs and DVDs; how to play streaming audio from Internet radio stations; and how to play and convert online audio files. When I buy a CD or DVD nowadays, I rip it and toss the original in a box in the attic. Who wants to juggle CDs and DVDs all day anyway? Finding my online audio and video archives is always exponentially easier for me than finding the CD that I accidentally used as a coaster in my office or the DVD that was accidentally buried under a stack of magazines. This sucks up a lot of disk space, but as far as I'm concerned, disk space is cheap compared to my time—in terms of both finding things and trying to keep them organized.

Chapter 20 explores one of everyone's favorite pastimes, playing computer games. Sad to say, most game manufacturers don't release their wares for Linux platforms. Luckily, this doesn't mean that you can't play games on Linux systems, just that you may not be able to play the latest and hottest 3D shooters, version of Grand Theft Auto, or other games that fans of the Spanish Inquisition like Jack Thompson are trying to ban. A huge variety of classic card and board games are preinstalled on Ubuntu and Kubuntu Linux desktop systems, as well as several clever (and fun) video games. Hundreds more are available on the Internet or are just a package manager session away.

Chapter 20

Would You Like to Play a Game?

Games on UNIX-like systems such as Linux have come a long way from "*Hunt the Wumpus*," a text-oriented adventure and exploration game that I spent many an hour playing in 1982 on a PDP-11 running primordial UNIX. Today, as you'd hope, Linux games are sophisticated graphical applications with plenty of performance and eye candy. One thing that you'll see in many of the games discussed in this chapter is that Linux games also have a sense of humor, which I always find refreshing. Losing and laughing is almost as good as winning (at least, that's what I tell myself).

As emphasized throughout this book, Ubuntu is a Linux distribution for people, and people are supposed to relax and have fun at times. Computer games are the aspect of desktop computer systems that give us all the opportunity to escape and relax without leaving our desks—hmmm, I'm not sure that this is exactly where I wanted to go when discussing relaxation, but *c'est le way it is*.

Ubuntu comes with a good complement of games preinstalled as part of your basic distribution. A huge number of other Linux games are available in the Ubuntu repositories and in source and package form on sites across the Internet. This chapter discusses the games that come with Ubuntu, highlighting the fact that you can get work done and have fun with the exact same machine. The first section of the chapter explores the games that come with a default Ubuntu system, the second section discusses the games that come with a default Kubuntu installation, and the last part of the chapter discusses classic games such as Chess, Go, Scrabble, and Monopoly, and how to play these (or similar games) on a Ubuntu or Kubuntu system.

IN THIS CHAPTER

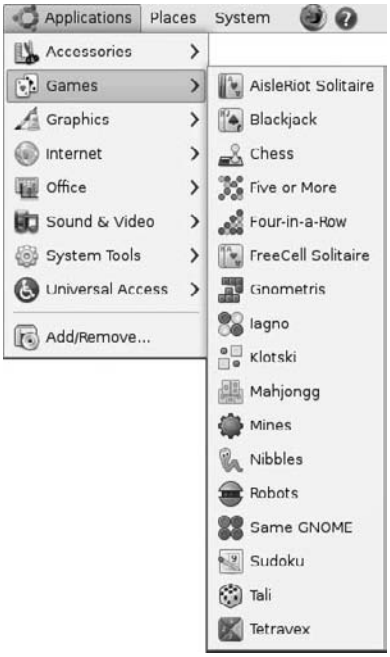
- Default games for Ubuntu
- Default games for Kubuntu
- Chess, Go, and Scrabble
- Arcade games

Default Games on Ubuntu

As shown in Figure 20-1, a default Ubuntu Linux desktop installation comes with a variety of different games across the gaming spectrum, which I classify as board games, card games, and video games.

FIGURE 20-1

The Games menu in a default Ubuntu installation



The rest of this section highlights the different games that come with a default Ubuntu desktop installation, providing a screenshot of each and an overview of the game. Because these all come with a default desktop installation at absolutely no cost and everyone enjoys different types of games, you'll quickly be able to figure out which games appeal to you and try them out yourself.

Card Games

Who hasn't walked past the desk of a Windows user sometime during the past 20 years and seen the quick flash as a Solitaire window was being minimized? Nowadays, this same surreptitious time sink is completely cross-platform, as Ubuntu Linux comes with some great card games. Once you master these, many more are available in the Ubuntu repositories. The following sections provide an overview of the card games that come with a default Ubuntu Linux desktop installation.

AisleRiot Solitaire

I've played solitaire since I was a kid, and somehow thought that there was only one solitaire game—that being, of course, the one that I'd learned to play, which turned out to be Klondike. I became more enlightened recently when my wife asked me what other versions of solitaire were available. Googling told me that there were indeed a few zillion solitaire games, so I bought her an analog data storage device that came preloaded with information about solitaire—that is, a book on the subject.

I was similarly enlightened when I suggested that she play AisleRiot Solitaire, which sports a mean game of Klondike and is installed on her Ubuntu-based laptop. Unlike me, she read the online help, which immediately told her that AisleRiot Solitaire actually provides 81 different solitaire games. (Ironically, this was a higher number than the ones discussed in the book I bought her.)

If you're a solitaire aficionado who knows more than one game, you'll be happy to know that AisleRiot Solitaire provides the following games: Agnes, Athena, Auld Lang Syne, Aunt Mary, Backbone, Bakers Dozen, Bakers Game, Beleaguered Castle, Block Ten, Bristol, Camelot, Canfield, Carpet, Chessboard, Clock, Cover, Cruel, Diamond Mine, Doublets, Eagle Wing, Easthaven, Eight Off, Elevator, Escalator, First Law, Fortress, Fortunes, Fourteen, Freecell, Gaps, Gay Gordons, Glenwood, Golf, Gypsy, Helsinki, Hopscotch, Isabel, Jamestown, Jumbo, Kansas, King Albert, King's Audience, Klondike, Labyrinth, Lady Jane, Maze, Monte Carlo, Neighbor, Odessa, Osmosis, Peek, Pileon, Plait, Poker, Quatorze, Royal East, Saratoga, Scorpion, Scuffle, Seahaven, Sir Tommy, Spider, One Suit Spider, Two Suit Spider, Spiderette, Straight Up, Streets and Alleys, Ten Across, Thieves, Thirteen, Thumb and Pouch, Treize, Triple Peaks, Union Square, Valentine, Westhaven, Whitehead, Will o' the Wisp, Yield, Yukon, and Zebra.

If you're like me and just love Klondike, it's in there, and it's a great implementation of the classic game.

Blackjack

Whether you're a blackjack lover and want to play a few thousand games without actually winning or losing money, or just want to practice your card-counting skills prior to a vacation in Vegas, Ubuntu's default desktop install comes with a nice online version of Blackjack that is a pretty standard multideck, casino-style blackjack game which also supports more advanced blackjack moves such as doubling down or splitting your hand. It uses Vegas Strip rules by default but can also be configured to use Atlantic City, Ameristar, or Vegas Downtown rules so that you can prepare for almost any gambling vacation.

Freecell Solitaire

Freecell is a fairly popular card game thanks to its appearance in various versions of Microsoft Windows over the years. For me, one of the nicest things about the GNOME games version of Freecell is that it provides a Hint button that is always handy as a last resort.

Implementation-wise, presenting this game as being standalone is a bit misleading, because this menu item just starts the AisleRiot Solitaire application but preloads a different card game.

Tali

Tali is an online version of a dice game that essentially emulates a card game (poker). If you're familiar with a popular analog game called Yahtzee from Hasbro, Inc., Tali may not be a complete shock to you, and may instead appear to be a familiar friend. By default, Tali is configured to enable one human player (you) to take on four computer-generated opponents, but you can easily configure your preferences to support multiple human players and reduce or eliminate the computer-generated opponents. I especially like the default prompt at the bottom of the screen when it's my turn: "Human!—You're up." Yes, Mr. Borg, sir!

Board Games

Board games have always been quite popular and don't have the sinful connotation that some people attach to card games. Although some of the games in this section are Linux implementations of existing computer games, I've defined board games as any game that uses a grid or could conceivably be played on a board.

Chess

Chess is one of the oldest and most popular games on the planet, so its no surprise that the Ubuntu folks provide a version of as one of their default games. I'm not going to explain the rules of chess here because there are thousands of books that are dedicated to doing that, plus it would quite simply take too long. The version of chess that is installed by default on Ubuntu systems is a version of GNU Chess, which comes with a very clever chess-playing engine that usually kicks my butt.

Beyond the game itself, one of the best things about chess on Ubuntu systems is the number of different versions of chess that are available. See the section later in this chapter entitled "Billions and Billions of Versions of Chess" for information about some of these, including those that enable you to play against real people online.

Five or More

Five or More is a GNOME version of a well-known Windows game called Color Lines. Like many good games, game play in Five or More is elegant in its simplicity. Each turn, three additional shapes appear on the screen in random locations. Your goal is to move existing shapes to try to arrange sequences of five matching discs, which removes them from the ever-more-crowded game field, somewhat like a slow-motion, nonlinear version of Tetris.

Unfortunately, Five or More is missing the Boss Alert key that was present in the Windows versions of this game, which was one of my favorite features. Pressing this key displayed a system utility screen that would make it look like you were doing real work while your boss was within range of your monitor. Fortunately, the game is just as addictive as Color Lines—you'll just have to hide the screen yourself when the boss approaches.

Iagno

Iagno is a disc-flipping game whose object is to flip as many of the other player's discs to your color, while protecting your own. You flip an opponent's disc by positioning two of your discs on either side of it. Iagno is a conceptual relative of the popular category of existing board games such as Othello and Reversi, and offers much the same sort of game play.

Klotski

Klotski, the game with one of the least pleasant names I've ever heard, is another game that many Windows escapees will be familiar with, although it has its roots in analog games from France, Thailand, and most directly from Poland, where "Klocki" is the word for blocks.

Klotski is a strategy game in which the object is to maneuver a large square out of the enclosed area by moving the smaller squares that surround it. It's quite fun to play and comes with a number of pattern variants, each of which tells you in advance how many moves should be required.

Mahjongg

GNOME Mahjongg is a single-player version of the classic Chinese Mahjongg game, which was apparently developed as a combination of earlier Chinese games. Mahjongg was extremely popular in the United States and Great Britain in the early twentieth century, and is still popular today. The object of the game is to remove tiles by matching pairs that display the same pattern. Mahjongg tile patterns are extremely intricate and can be quite beautiful, as they are in the graphics for this version of the game. The GNOME Mahjongg game is quite fun and addictive, and also features a nice Hint button for use when your eyes blur after staring at too many intricately patterned, but very similar, tiles.

Mines

Mines is the GNOME version of the popular Microsoft Windows Minesweeper game. Clicking on a square reveals the number of mines in area around that square. Once you identify a square as containing a mine, you mark it by planting a flag on that square. If you clear or mark all squares on the board without clicking on a mine, you win. You are then judged on the time that it took for you to reach the end.

I've never particularly understood why so many people find this game interesting, but I feel exactly the same way about the GNOME and Windows versions, so I suppose that this means that the GNOME version is quite faithful to the original.

Sudoku

If you aren't familiar with Sudoku, I envy you the rock that you must have been hiding under for the last few years. Hundreds of books of Sudoku games are available, many online sites provide daily Sudoku games, and I seem to see a new one of these every day in many of those daily paper things—oh, right, newspapers.

Sudoku is a number placement game, where the object of the game is to fill the game grid such that each column, each row, and each of the 3×3 block subsets of the game field contains each of the digits from 1 to 9 only once. Starting a game of Sudoku salts the game field with a starting set of numbers, and you supply the rest. Sudoku is incredibly popular and, although the craze seems to have peaked somewhat, it is still as sublimely addictive as any great game.

Tetravex

Tetravex is an interesting game where you place squares with numeric values beside existing squares where the values on each edge match. I find Tetravex to be quite fun—and like the kindest GNOME games, its File menu provides a Hint key that hastens the learning process and lets you win early on, while you're still evolving your strategy.

Video Games

Video games have probably gotten more kids interested in computers than any other motivator. After all, it's not surprising to see a video game in a bar or grow up staring at your PlayStation and find yourself thinking, "I wish I could create something like that." This is a far more likely scenario than staring at a huge report in your father's or mother's briefcase and thinking, "I wish I could write a program that would produce neat reports like that!"

The GNOME games pack installed on Ubuntu Linux desktop systems has some great implementations and different takes on some classic video games, and many well-known games such as Quake are directly supported under Linux. For information about actually playing actual arcade classics on Linux using images of the original ROMs, see the information on MAME (Multi-Arcade Machine Emulator) in the "Popular Emulation Software" section of Bonus Chapter 7.

Four-in-a-Row

Four-in-a-Row is an interesting and addictive cross between tic-tac-toe and Tetris, which should help keep you from doing real work for many hours. Your goal in Four-in-a-Row is to be the first to arrange a row of four pieces of your color, while preventing your opponent (or computer opponent) from doing so. You click in a row to drop and stack a piece in that row, adding to your sequence or blocking your opponent's.

Gnometris

Gnometris is a GNOME implementation of Tetris, probably the most famous computer game of all time after Pong. Invented in Russia, it has all of the elements of a great game—it's simple to grasp, easy to play, and when you lose you always find yourself thinking, "I was so close—I can beat this next time!"

In case you grew up in a cave or under a rock, the object of all Tetris-like games is to orient falling shapes so that they complete rows, which you then score points for as they disappear. The pace of the game accelerates as time passes, and once a shape reaches the top row of the screen without completing a row, the game is over. Simple, fun, fast, and addictive.

Nibbles

Nibbles is a simple but classic game in which the user controls the movements of a snake that travels around the screen, eating diamonds of various colors and growing in length as the game progresses.

Each player has three lives, one of which is used up each time the snake hits a wall or itself. If you're playing by yourself, the game ends when you run out of lives.

Nibbles is interesting because it is also network-aware. The Game ➪ Network game menu command displays a dialog that enables you to identify and connect to a games server that enables you to play Nibbles against other users who also have access to that server. As you'll see later in this chapter, many of the classic board games such as Chess and Go also feature networked modes, but Nibbles is a simple and fun example of a game that you and your friends can play together over a network, each from the comfort of your own console.

Robots

GNOME Robots is a graphical version of the many text-based hunt-and-kill games available on older UNIX systems. All of the robots in the game want to kill you, and your motivation is simple. First of all, you must survive, and second, you must trick the robots into destroying themselves, either by colliding with each other or by colliding with the wreckage left behind by a previous collision. This game can be surprisingly addictive.

Same Gnome

Same Gnome is a puzzle game that combines aspects of Tetris, Five or More, and similar games, in which you remove areas of similarly colored pieces on a grid by clicking on them, which causes the pieces around them to drop down or over, hopefully increasing the size of other similarly colored areas which you can then remove, and so on.

This is really a fun little game that can be surprisingly addictive. It's so simple that you'll continue to think, "I can beat this!" long after every other living creature in your house or apartment has fallen asleep.

Games on Kubuntu Systems

A tremendous number of free games are available for KDE-based distributions such as Kubuntu. Unlike GNOME-based Ubuntu systems, games are not installed on Kubuntu systems. To install them, you must use your favorite package management utility to install the `kdegames` package on KDE 3-based Kubuntu systems (i.e., Kubuntu 8.04 and earlier), or the `kdegames-kde4` package on KDE 4-based Kubuntu systems (i.e., Kubuntu 8.10 and later). Installing these packages requires less than 50 MB of free disk space, so it's not as though installing them will bring your system to its knees (playing them all the time might, however).

TIP **If the KDE games discussed in this section aren't enough for you and you're intrigued by any of the games discussed earlier in this chapter, you can always install the `gnome-games` package on your Kubuntu system to help ensure that you'll rarely get any real work done.**

Unlike the Applications ➪ Games menu on GNOME-based Ubuntu systems, Kubuntu's K Menu ➪ Games menu (K Menu ➪ Applications ➪ Games on KDE 4-based Kubuntu systems) is nicely organized into different classes of games, each of which (and the games that it contains) is discussed in the next few sections. Figure 20-2 shows the Games menu on a KDE 3-based Kubuntu 8.04 system.

FIGURE 20-2

The Games menu in a default Kubuntu 8.04 installation

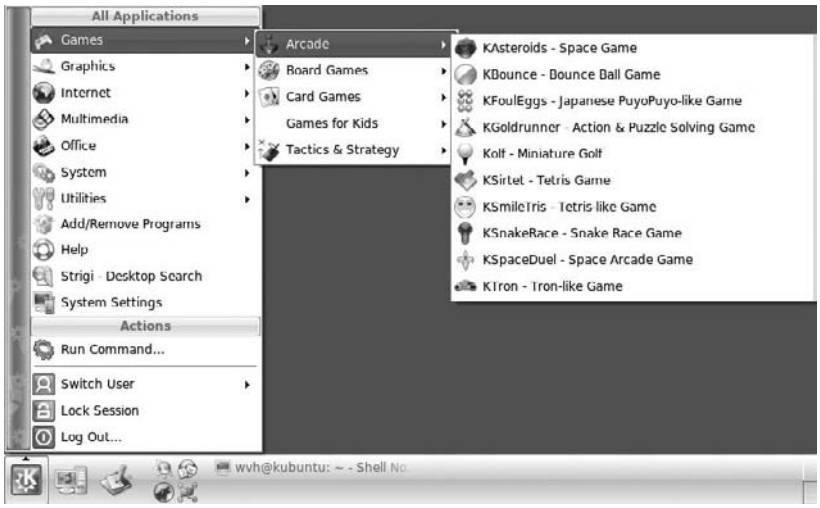


Figure 20-3 shows the Games menu on a KDE 4-based Kubuntu

The following sections are organized by the name of the game rather than by their description, so they are not necessarily discussed in the same order that you'll find them on the appropriate Kubuntu game menus. Games that are only provided by default on KDE 3-based or KDE 4-based Kubuntu systems are identified as such—the descriptions of games that do not specify any version information are provided by default on both types of Kubuntu systems (although, of course, they look slightly different). Games that are not provided by default on either type of Kubuntu system may still be available in the Ubuntu repositories.

Arcade Games

Arcade games are the things that first got me interested in computer science, because someone explained that they had a computer inside that controlled and displayed the action. I might not have had to take student loans if I still had all of the quarters that I invested in the games themselves, but things have still worked out OK.

FIGURE 20-3

The Games menu in a default Kubuntu 8.10 (or newer) installation



The `kdegames` package provides a number of excellent versions of many of the most popular arcade games from "back in the day." Even if you weren't born in time to actually see them, these games became classics by combining intriguing story lines, good game play, and the ability to make players believe that they could have won were it not for one mistake or bit of randomness. The KDE implementations of these games preserve everything about these classic games except for a roll of quarters.

KAsteroids

Only available on KDE 3-based systems, KAsteroids is a faithful adaptation of Asteroids, one of the early arcade games. The object of the game is to blast passing asteroids with the nose cannon mounted in your spaceship. As you shoot an asteroid, it breaks into multiple smaller pieces, which you must also shoot, and which also break into multiple smaller pieces. Your controls are somewhat minimal—you can initially spin your ship, fire its rockets, and fire the nose cannon. You have a limited amount of fuel, but you can occasionally pick up more by flying over a fuel symbol. You can also upgrade your ship, which allows you to hit the brakes and provides shields.

If you remember the old Asteroids arcade game, KAsteroids is a nice translation of the traditional vector graphics game into a raster graphics environment, with a few extra bells and whistles to while the night away.

KBlocks

Only available on KDE 4-based systems, KBlocks is a nice implementation of Tetris—no special features, no new rules, just a nice-looking version of Tetris.

Like most Tetris games, the game board displays an area in which blocks drift to the bottom of the board, and can be rotated, moved left to right, and accelerated when you're sure that they will fall in the right location. Fifty million Tetris fans can't be wrong—and if you like Tetris, you'll enjoy this game.

KBounce

KBounce is an interesting game that is reminiscent of arcade games such as Qix, where the object of the game is to wall off a majority of the playing field from the moving elements within it. In KBounce, the moving objects are initially two small balls, although a new ball is added each time you move to a new level by walling off more than 75% of the existing playing field.

To wall off a portion of the screen, click the left mouse button on any square, at which point a wall will begin to spread in both directions from that square. Whether the walls are horizontal or vertical is toggled by clicking the right mouse button.

KBreakout

Only available on KDE 4-based Kubuntu systems, KBreakout is a great version of the classic Breakout game, a linear descendent of the early Pong tennis game. The object of the game is to remove bricks from a wall by bouncing a ball off them. When the ball hits a brick, that block is removed and the ball bounces back towards your paddle. The real fun begins if you manage to bounce a ball through an empty column so that it bounces around behind the wall, eliminating bricks with no effort on your part until it emerges through an empty column.

KBreakout adds some nice twists to the Breakout concept. Each level features a different wall of bricks, some with predefined empty columns or immutable bricks. Eliminating some bricks drops prizes that you can catch to get extra points or special capabilities, such as a double-wide paddle, an extra ball in play at the same time, and so on. KBreakout is long-term fun for the whole family.

KFoulEggs

Only available on KDE 3-based systems, KFoulEggs is a game in which you attempt to align falling, multicolored tiles so that multiple portions of the same color are adjacent. Whenever you maneuver two or more tile segments of the same color beside each other, they join together. Once four or more are joined, that segment disappears.

According to the documentation, KFoulEggs is an implementation of a popular Japanese arcade game known as PuyoPuyo. This was useful information, because otherwise I would have thought that it was just yet another “Tetris with a twist” game.

KGoldrunner

KGoldrunner is a maze and ladders game that is evocative of the old personal computer Lode Runner game. Pursued by enemies, you navigate through each level by running and climbing ladders, picking up gold to score points, while digging pits to trap your pursuers.

KGoldrunner comes with multiple levels with varying degrees of difficulty. One of its nicest features is that it comes with a level editor that you can use to create your own levels. (Lode Runner was one of the games that introduced the idea of user-created levels.) Being able to expand the game ensures that you’ll never get bored with it, but can also guarantee that you’ll never be done with it.

Kolf

Kolf is a KDE implementation of miniature golf, which I think is a much more interesting game than trying to emulate the actual game of golf online. You orient your putter with the mouse, and then hold down the left mouse button for varying periods of time to control the intensity of your stroke.

Kolf combines some of the aspects of real golf (such as water hazards) with the types of action that you traditionally find on miniature golf courses (tunnels, walls to bounce the ball off of, and so on). It is also

faithful to another drawback of many miniature golf courses, which is that it does not provide a 19th hole. Your refrigerator will have to do.

Kollision

Only available on KDE 4-based systems, Kollision is an interesting game in which the object of the game is to move the ball that you control so that it avoids being struck by the other balls that bounce around the edge of the screen. The longer you avoid other balls, the more points you get, but additional balls are added to the screen at regular intervals, increasing the number of obstacles.

As you might expect, Kollision is quite easy initially, which makes it easy to get sucked into as the number of balls that you must avoid increases. Bouncing balls change their angles not just by hitting the wall but also by hitting each other, which makes the game a great way of studying Brownian movement while having some fun.

KSirtet

Only available on KDE 3-based systems, KSirtet is a KDE Tetris game that looks remarkably like KFoulEggs, sharing almost everything except for the appearance of the falling tiles. It was also developed by the same individual, Nicolas Hadacek, demonstrating the power of code reuse.

In case anyone is wondering, KSirtet is Tetris spelled backwards with the traditional K glued on the front. Perhaps I'm the only person who didn't spot that immediately.

KSmileTris

Only available on KDE 3-based systems, KSmileTris initially appears to be simply yet another Tetris clone that happens to use multiple types of smiley faces as the components of its tiles. Actually playing the game reveals that there are many other differences, including how tiles can change as they drop. In most Tetris clones, tiles simply rotate as they drop—in KSmileTris, you must align both the types of smiley faces and the tiles themselves.

Having to orient different types of smiley faces, as well as the geometry on which they are displayed, is trickier than it sounds. KSmileTris is a nice twist on Tetris, and lots of fun.

KSnakeRace

Only available on KDE 3-based systems, KSnakeRace makes you a hungry snake that must navigate the screen to eat all of the apples with a set period of time, while avoiding a computer-controlled snake, the edges of the screen, and various obstacles.

KSnakeRace is fast-paced and surprisingly difficult. It enables you to create and use your own levels, although this involves editing bitmaps of the default screen. That may seem like a lot of hassle, but was the only way that I could beat the game on a fast, modern system—I think that my high score on the default level is 11.

KSpaceDuel

KSpaceDuel is a two-player (or one player against your computer) game in which two ships fly around a sun and try to destroy each other by shooting at each other or laying mines. Keyboard commands enable you to rotate your ship, accelerate, lay mines, and so on.

KSpaceDuel is evocative of the original computer video game, SpaceWar, but requires KDE rather than a PDP-1. KSpaceDuel is also more complex than it looks—for example, the sun is not just an obstacle but has a gravitational field that affects how ship, bullets, and mines behave. Mastering the keyboard controls took me quite a while, however— but once past that, the game itself is fun.

KTron

Only available on KDE 3-based systems, KTron is a simple version of the motorcycle racing game from the classic camp movie "Tron." By simple, I mean that you (and another player, or the computer) each control a single square and move it around the screen, avoiding the other player and the edges of the screen. The object of the game is twofold—don't hit anything yourself, and try to box in your opponent to ensure that he/she does.

The fact that KTron doesn't have the fancy racing motorcycles or the 3D viewpoint of many other KTron clones means that it is basically a standard snake-type game without many bells and whistles. Devotees of that type of game may enjoy it more than I do.

Board Games

Board games are defined as games on which game pieces are placed, moved, and removed on a premarked game board. Identifiable board games date back to the Bronze Age, 3500 BCE or so, and were probably around even earlier. The identity of the first person to suspect that "Centurion Mustard did it in the great pyramid with a ceremonial dagger" is lost in time. Luckily, no game companies from this era are still active, or they'd probably still hold the copyrights.

Almost everyone has played board games at one time or another, and the `kdegames` package includes a number of unique board games as well as some that were inspired by commercial board games.

Atlantik

Only available for KDE 3-based systems, Atlantik is a KDE client for the Monopoly-like Atlantik game that enables you to play the networked version of Atlantik against remote opponents.

If your Internet connection is down or you simply can't afford the additional external bandwidth, you can install and configure the server package (`monopod`) on your own system so that you and your siblings, spouse, and friends can at least play each other. See the section later in this chapter entitled "Monopoly Clones" for more information and some actual screenshots of the game, which is great once you actually locate opponents.

KBackgammon

Only available on KDE 3-based systems, KBackgammon is a nice-looking version of backgammon that enables you to play against your computer, other players, and networked opponents via FIBS (the First Internet Backgammon Server). It also provides experimental support for backgammon engines such as `gnubg` (GNU backgammon, `www.gnubg.org`), but that doesn't actually seem to be implemented.

KBackgammon is great if you want to play another person or participate in the online backgammon community.

KBlackBox

KBlackBox is an intriguing logic puzzle game where you must identify the location of four balls that are hidden on the board. Your only guidance for the location of the hidden balls is the ability to shoot rays into the box by clicking on balls on the border of the playing field and seeing where they emerge, hopefully after bouncing off one or more of the hidden balls. After placing your four balls where you think the hidden ones are located, you click the Done icon to see how you fared. The hidden balls are revealed and your final score is displayed.

Your score is computed by the correctness of your guess and the number of rays that you shot while trying to locate the hidden balls. Once you master the rules and figure out how the rays work, this is a great game for people who enjoy logic puzzles.

KBovo

Only available on KDE 4-based systems, KBovo is an interesting Tic-Tac-Toe variant played in a large grid. The object of the game is to get five squares in a row, either linearly or diagonally. Your opponent is always the computer.

KBovo's larger playing field and graffiti-like graphics makes game play feel a bit like you're designing plays for your local football team, but the core game is fast and fun.

Kenolaba

Only available on KDE 3-based systems, Kenolaba is a two-player strategy game that is very similar to the well-known European board game Abalone. Kenolaba is Abalone spelled backwards, with the traditional K glued on the front.

Both you and the computer begin with 14 marbles, and the object of the game is to push the other players' marbles off the board. You can move one, two, or three marbles inline or diagonally, one space at a time. You can push the opposing player's marbles a single space in an inline direction when you have numerical superiority in that row.

Kenolaba is a fast-paced and enjoyable game that is easy to grasp, and similarly easy to spend hours playing.

KFourInLine

Only available on KDE 4-based systems, KFourInLine is essentially the same game as KWin4 on KDE 3-based systems. The game "board" consists of seven columns of six spaces each. Pieces are added from the bottom of the column up, like a coin sorter. The object of the game is to position four of your pieces in a row, either linearly or diagonally.

Like KWin4, KFourInLine is simple and addictive. Having to think a few moves ahead when placing pieces to win as well as blocking your computerized opponent can be quite a challenge—at least, apparently, for me.

Kiriki

Only available on KDE 4-based systems, Kiriki is a Yahtzee-like game that is designed to be played by up to six players. You collect points by rolling up to five dice up to three times per turn, trying to match predefined sets of score patterns.

By default, you are always the first player, and the other five players are computerized opponents. A game consists of 13 rounds.

KMahjongg

KMahjongg is a single-player version of the classic Chinese Mahjongg game, which was apparently developed as a combination of earlier Chinese games.

Mahjongg was extremely popular in the United States and Great Britain in the early twentieth century, and is still popular today. The object of the game is to remove tiles by matching pairs that display the same pattern. Mahjongg tile patterns are extremely intricate and can be quite beautiful, as they are in the graphics

for this version of the game. KMahjongg game is quite fun and addictive, and also features a nice Hint button for use when your eyes blur after staring at too many intricately patterned, but very similar, tiles.

KReversi

Only available on KDE 4-based systems, KReversi is a great implementation of the classic Reversi and Othello disc-flipping games. The object of the game is to flip as many of the other player's discs to your color as possible, while protecting your own. You flip an opponent's disc by positioning two of your discs on either side of it.

KReversi's game play engine provides hours of solo fun, as well as great practice if you ever play against another human.

KSquares

Only available on KDE 4-based systems, KSquares pits you against human or computer opponents, where the object of the game is to draw lines segments between points on a grid, eventually enclosing areas that become your property.

Like many other "define your territory" games, KSquares is simple and addictive. Initially, it's a race to place line segments that do not define spaces, but once those disappear, it turns into a "who can sacrifice the least territory" game, which is lots of fun.

KWin4

Only available on KDE 3-based systems, KWin4 sounds like a yet another KDE window manager but is actually a two-player game played against the computer or a local or networked opponent. The game "board" consists of seven columns of six spaces each. Pieces are added from the bottom of the column up, like a coin sorter. The object of the game is to position four of your pieces in a row, either linearly or diagonally.

KWin4 is simple, fun, colorful, and addictive. The ability to play against the computer is perfect for both learning the game and passing an hour or two of free time.

Shisen-Sho

Shisen-Sho is a Mahjongg-like game that is played with the same tiles as Mahjongg. Instead of stacking tiles, the tiles are laid out on a single 12×24 grid. Matching tiles can only be removed if they can be connected with a maximum of three connected lines.

Shisen-Sho is fun and challenging. The size of the grid can make it difficult to spot tiles that can be removed—in which case, pressing the "H" key displays a recommended move. Love those hints. . . . A game of Shisen-Sho is over when no moves are possible, and some games are unsolvable (you can eliminate unsolvable games using Shisen-Sho's configuration dialog). The final score is calculated using an equation that makes my brain hurt just to look at. For me, the game is more fun when I think of it as a form of solitaire, not as a scored game.

Card Games

Card games have been around since the tenth century, and are just as popular today as they probably were back then. A deck of cards provides cheap, portable fun almost anywhere. Even if you still insist on ordering "freedom fries," you have to thank the French for playing cards as we know them today, with the exception of the joker and strip poker, both of which were invented by Americans.

KPoker

Poker is fun—just ask the zillions of people who now watch it on TV and play it over the Internet. Doesn't anyone still play with actual cards in smoke-filled back rooms anymore? Only available on KDE 3-based Kubuntu systems, KPoker makes a great opponent if you want to hone your skills without actually losing (or winning money). The only thing missing is a virtual cigar to smoke after winning (and, of course, real money).

By default, KPoker is just you against the computer, but you can have the computer play multiple hands to create a bit more virtual excitement. Bluffing isn't as much fun as it is with real poker, since if my computer has a tell, I can't spot it.

Lieutenant Skat

Maybe I don't play cards enough, but I originally thought that Lieutenant Skat would be a singing game in which I was up against Ella Fitzgerald in a WAC/WAV uniform. Surprisingly, that's not the case at all—instead, Lieutenant Skat is a two-player trump/war game that is based on an existing German game called Offiziersskat. Like its inspiration, Lieutenant Skat is played according to a simplified version of the rules of the traditional Skat card game.

You can play the game against the computer, against another person, or over a network, though I haven't tried the networked version. Lieutenant Skat is fast and fun—and someday I will beat the computer.

Patience

Patience is the British version of Solitaire. As I mentioned in the section on Ubuntu's AisleRiot Solitaire game earlier in this chapter, I've played Solitaire since I was a kid, but never realized that there were actually hundreds of different Solitaire games. The version I always played turned out to be called Klondike, which is the default game that is dealt by Kubuntu's Patience.

Like Ubuntu's AisleRiot Solitaire, Kubuntu's Patience supports many different Solitaire games. The intersection between the ones that it supports and the ones that are supported by Ubuntu's AisleRiot Solitaire is surprisingly small. If you're a Solitaire, er, Patience fan, Patience is another good argument for mixing GNOME and KDE applications on a single system—unless your goal was actually to get some work done on your computer system.

Games for Kids

Although many Linux games have been designed for children, the `kdegames` package only includes one example. Don't worry—plenty of others are available!

Potato Guy

Potato Guy, also known as KTuberling in the grand KDE tradition of names beginning with K that avoid copyright issues, is intended for small children and is not so much a game as a toy. Like the physical Mr. Potato Head game, the object of the game is to amuse yourself by putting eyes, a nose, a mouth, and various hats and badges on a potato. As you'd expect, the potato in this case is online, which is both cost-effective and removes the chance of finding a withered potato under your couch or in your child's toy box one day.

By default, Potato Guy uses a speech mode that verbalizes the name of each piece that you position. If you install multiple languages on your system, your kids can learn words like eye, mouth, nose, and so on in those languages. Not exactly a Berlitz course, but still useful. Potato Guy doesn't limit you to just adding

features to a potato—as a Linux game, it also allows you to add them to Tux, the Linux Mascot. The game also allows you to position fish and other undersea objects in a virtual aquarium, which has a yawn factor of several thousand. Once you've completed any of these, you can save them as a picture in various formats, which your child might enjoy using on their desktop.

Other Games for Kids

For suggestions about other games that your kids may find fun (and that may also be educational), the following Web sites provide some great summaries, suggestions, and links:

- `www.layhands.com/FreeGamesForKids/Linux.htm`
- `www.fanaticattack.com/2008/exceptional-linux-programs-for-kids.html`
- `www.makeuseof.com/tag/muo-games-10-open-source-online-games-for-kids`
- `www.happypenguin.org/list?sort=category&search=kids`

Tactics & Strategy

Kubuntu's strategy games typically provide a nice combination of video action and fun game play. As the name suggests, all strategy games require some planning, forcing you to think ahead a few moves—both to where you want your game pieces to be, and where you think that your opponent is going.

KAtomic

KAtomic is a Sokoban-like game (see the discussion of KSokoban later in this section) in which the object is to create molecules by attaching atoms of various sorts—hence the name. The object of the game is to do so in as few moves as possible.

Like Sokoban, KAtomic is more complex than it initially seems. Selected molecules move until they hit an obstacle, which means that you must first decide where all of the atoms should eventually congregate and then figure out how to get them there. KAtomic is a great way to accidentally learn something (the atomic composition of various molecules) while having a great time.

KBattleship

KBattleship is a great implementation of the classic Battleship game with cute graphics and great sound effects. You can play KBattleship against the computer, a local opponent, or over a network. It's great fun, though networked game play doesn't allow you to hear your opponent's anguished cry of "You sank my battleship!"

KJumpingCube

KJumpingCube is a fun strategy game in which the eventual object is to own all of the squares on the board. Each square has both an owner and a point value. In each of your moves, you can either grab a new square or add points to existing ones that you own.

The key element of seizing adjacent (linear) cubes is that if a square has more points than it has neighbors, the points jump to its neighbors—and take them over. This can cause some amusing ripple effects in game play.

KJumpingCube is a fun game that requires that you think about the board layout, the point scores in cubes all over the board, and whether to stake out new territory or pump up the point values in some of your squares for the future. Lots of fun!

Klickety

Only available on KDE 3-based systems, Klickety is an extremely clever version of Tetris in reverse. The game field starts out completely populated with squares of various colors. You click on specific squares to remove them and all adjacent squares of that color. The squares above them then slide down, filling the empty spaces that you've just created.

I have lost more hours to Tetris than I care to admit in print, and this game is equally fun. Kiss a few hours goodbye.

KMines

No package of computer games would be complete without a version of Minesweeper, the Microsoft Windows game that is probably responsible for almost as much lost productivity as blue screens and subsequent reboots. The `kdegames` package delivers KMines, which is a faithful adaptation of Minesweeper.

KMines provides multiple levels, and also enables you to create custom levels with a configurable percentage of mines. Personally, I've never been a big fan of Minesweeper and I don't really enjoy KMines, so I guess that means it's a pretty faithful adaptation.

KNetwalk

KNetwalk is a fun game for nerds. The object of the game is to rotate segments of network cabling and computer connections so that every system is connected to a server, and no network segments are unconnected. You rotate segments and computer connections by clicking on them, and your score depends on both success and the number of clicks that it takes to connect everything.

KNetwalk is so faithful to the computer field that there is little documentation for it. Although the concept is fairly simple, the game is actually lots of fun—and you may be able to convince your boss that you're doing network planning if he or she spots you playing it.

Kolor Lines

Kolor Lines is a single-player game where you're essentially playing against the existing high score, and was inspired by Color Lines, a popular PC game of long ago. To score points, you must position five balls of a single color side by side in a row or column, at which point you receive points and they disappear. You can move any ball to any square that it can reach by any number of horizontal or vertical moves. As you move each ball, the computer adds three more to the screen.

This is another game in the "I almost won that time" school of addictive game design. The computer's placement of the new balls adds randomness that can easily work for or against you, but will certainly keep you playing for longer than you'd planned to.

Konquest

Konquest is a strategy game with the conservative goal of galactic domination. You can play against any number of computer players, and create different numbers of neutral planets as fodder for your empire building.

Each planet can produce a certain number of ships during a turn, and you attach other planets by selecting one of your planets as the source, the world that you want to conquer as your destination, and identifying the number of ships that you want to commit. You can get information about the status of your opponents' planets and their combined fleet strength at any time.

You may find that it takes some time to understand all of the rules, the interaction between planets and opponents, and how to launch a successful campaign, but then Konquest almost always takes a while.

KSokoban

KSokokoban is a nice version of KSokoban, the Japanese warehouse manager game in which the object was to position identical crates beside each other. KSokoban is a bit more fanciful, and your character has to position gems on the green goal squares on the screen. To move your player, left-click on any square that can be reached from your current location without being blocked by a gem or a wall. Once your character is beside a gem, right-click on the gem to push it one square forward.

Planning how to move gems so that you can subsequently still push them while not blocking yourself from doing so can be quite challenging. KSokoban includes five different collections of levels, including easy ones that you can practice on until you get the hang of the game.

KSameGame

KSameGame is a single-player game that was inspired by the popular Mac OS game called SameGame, itself inspired by earlier games in which the object is to remove all of the marbles on the screen by clicking on groups of adjacent marbles of the same color. When you remove marbles, the remaining ones drop down, and all pieces to the left of the removed marbles move to the left.

A game of KSameGame ends when there are no more marbles that can be erased, whether because you've cleared the screen or (more likely) because there are no more adjacent marbles of the same color. Like many other great games, KSameGame (and its predecessors) provides an addictive combination of random marble placement and strategy.

KSirk

Only available on KDE 4-based systems, KSirk is inspired by and very reminiscent of the Risk board game. Each player controls armies that can be positioned in countries that you own around the world. You can attack adjacent countries by selecting a country until it turns red, then selecting the country that you want to attack.

Though the sound effects can be a bit cheesy, they're much better than those in the Risk board game (which is silent), and the computer does a good job of injecting some randomness into your attacks on other countries. If you ever wanted to rule the world but didn't want to leave your office chair, KSirk is the way to do it.

Games on the Games Menu

If you are running Kubuntu 8.10 or later, or have installed the KDE 4-based Kubuntu 8.10 games pack on your system, the games menu also includes three games that are not on any submenu. I'm not sure whether this is for convenience because of their popularity or because they don't exactly fit any exiting submenu category, but whatever the reason, these games are all fun to play.

KDiamond

Only available on KDE 4-based systems, KDiamond is an interesting game that initially displays a 10×10 grid filled with gems of various types. The object of the game is to swap the positions of any two adjacent gems so that you create three gems of any type in a horizontal or vertical row. These three gems then disappear along with any other adjacent gems of the same type, you get points, and the gems above them drop down to fill the empty slots.

Each game is timed at 3 minutes and 20 seconds, so you must remove as many gems as possible within that period of time. This is a fun game, with enough randomness in the new gems that appear from the top of the screen to keep you playing for far longer than you anticipated.

KSudoku

As mentioned when discussing the GNOME version of Sudoku, earlier in this chapter, I envy you the rock that you must have been hiding under for the last few years if you aren't familiar with Sudoku. Hundreds of books of Sudoku games are available, many online sites provide daily Sudoku games, and many local newspapers include daily syndicated Sudoku games. Sudoku is a number placement game, where the object of the game is to fill the game grid such that each column, each row, and each of the 3 × 3 block subsets of the game field contains each of the digits from 1 to 9 only once.

Starting a game of Sudoku salts the game field with a starting set of numbers, and you supply the rest. Sudoku is incredibly popular and, though the craze seems to have peaked somewhat, is still as sublimely addictive as any great game.

Kubrick

Only available on KDE 4-based systems, Kubrick is a nice online implementation of a Rubick's Cube. You can display 3-dimensional views of the cube from the front, the cube from the front and back, or a 3D view of your cube with front and back views off to the side. The multi-cube views are especially useful in determining exactly what will happen when you rotate a given row.

The default Rubick's cubes that are provided in the game are 3 × 3 × 3 cubes. You can select cubes of other sizes, or even create your own, by selecting other difficulty levels and specific cubes from the sub-menus on the Game ➪ Choose Puzzle Type menu. Kubrick supports cubes up to 6 × 6 × 6 by default, and also includes some non-cubic forms, such as 6 × 3 × 2 and 6 × 4 × 1 shapes.

Other Popular Games

The games packs installed as part of default Ubuntu or Kubuntu desktop installations are really just the tip of the iceberg as far as games on Ubuntu and Kubuntu Linux is concerned. The Ubuntu repositories each offer many other games that are a Package Manager session away. Trying to describe each of these games would bloat this book far beyond its purpose of helping you to get up and running with all aspects of Ubuntu and Kubuntu Linux, so this section focuses on how to find equivalents or clones of some well-known games that many people would enjoy playing on their systems if they only knew what they were called and how to use them. For information on installing packages, see Chapter 23, "Adding, Removing, and Updating Software." The remainder of this chapter focuses on explaining what to look for if you're interested in a specific game or type of game.

Billions and Billions of Versions of Chess

In the world of open source computing, chess has two things going for it: it's incredibly popular, and nobody owns it. The fact that chess and its ancestors have been popular for a few thousand years as a thought-provoking strategy game makes it no surprise that versions of computer chess have existed since computers had enough memory to manage any sort of sophisticated game play. The fact that nobody owns the rights to the name "chess" means that there are no royalty or licensing entanglements, so any computer chess game can actually be called chess without having to resort to camouflage names or slight changes to the rules to escape corporate greed.

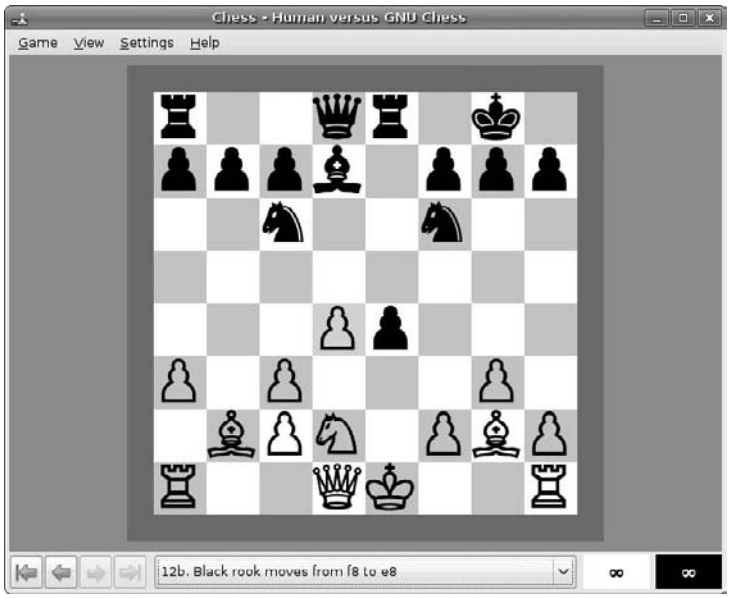
A simple search for chess against package names and descriptions in the Synaptic Package Manager shows 26 packages. These include local and network-enabled chess games, chess-playing engines that you can compete against, a chess game that you can play inside Emacs, special fonts for typesetting chess moves and even boards that you can format using TeX and LaTeX (discussed in Chapter 16), and much, much more.

TIP **If you're a true chess fan, your competition is no longer limited to friends, family, and neighbors thanks to the Internet. You can find several chess servers on the Internet, through which you can play remote opponents from the convenience of your own workstation. For more information, see Web sites like the Free Internet Chess Server site at `www.freechess.org` and Tim Mann's great summary of worldwide chess servers at `www.tim-mann.org/ics.html`. There are also a variety of chess servers that require some sort of payment or subscription to use their services and play against their subscribers.**

Figure 20-4 shows the version of chess that is installed by default on Ubuntu desktop systems. This is GNU Chess, which is a version of chess from the GNU folks that also includes a great chess-playing engine that enables you to play against your computer long after all of the other humans in your house have turned you down or gone to sleep. The GNU chess engine is probably the best-known local chess engine for Linux (and most) systems, because it is GNU software and therefore free.

FIGURE 20-4

The GNU chess interface



If you get angry at the GNU chess engine for always kicking your butt or simply want to experiment with others, the alternate fruit, phalanx, and sjeng chess engines are also available in the Ubuntu repositories. Unfortunately, the default Ubuntu version of chess can't be configured to use these other chess-playing engines (it automatically uses its internal GNU Chess engine), so you'll have to install another chess game that is a bit more social and therefore supports different chess engines. On Ubuntu and Kubuntu systems, a good candidate is `xboard`, which is a generic interface for playing chess against a variety of chess engines. Installing the xboard package adds an entry for this game to the standard Applications ➪ Games menu.

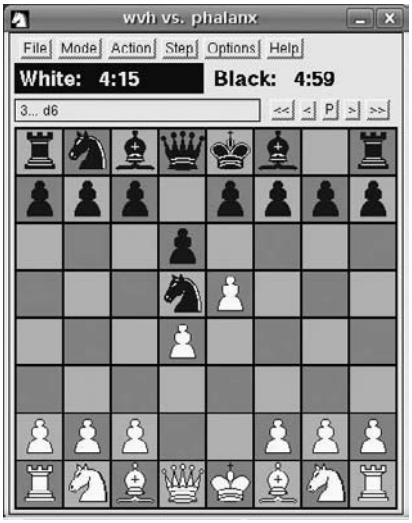
After installing `xboard` and the other chess engines that you want to try, you start `xboard` from a command line and use the `-fcp` (first chess program) command-line option to specify the chess engine that you want to play against. For example, the following command starts `xboard` using the `phalanx` chess-playing engine:

```
/usr/games/xboard -fcp /usr/games/phalanx &
```

Figure 20-5 shows the screen that you would see after executing this command.

FIGURE 20-5

Using xboard and a specific chess engine



As you can see from Figure 20-5, the header identifies the players—in this case, me (wvh) against the `phalanx` chess engine. Playing against another chess engine simply requires that you specify the path to the other chess engine as the argument to the `-fcp` option. If you want to play an Internet opponent, you can use the `-ics` option to specify an Internet Chess Engine to play against. (The `-fcp` and `-ics` options are mutually exclusive—you can only play one chess entity at a time.)

TIP **Though the default version of chess installed on Ubuntu systems is GNU Chess, the GNU chess engine that applications like `xboard` must use (`gnuchessx`) is not installed by default. To use `xboard` with the GNU chess engine, you should install the `gnuchess` and `gnuchess-book` packages using your favorite package management software, as described in Chapter 23.**

Unless your screen is much larger than mine, you may want to use `xboard`'s `-size` option to specify the size of the game board that it displays when it starts. I typically use `-size small` when playing `xboard`, but you may prefer other sizes such as Medium, Moderate, or Average—see the main page for `xboard` for a complete list of available sizes.

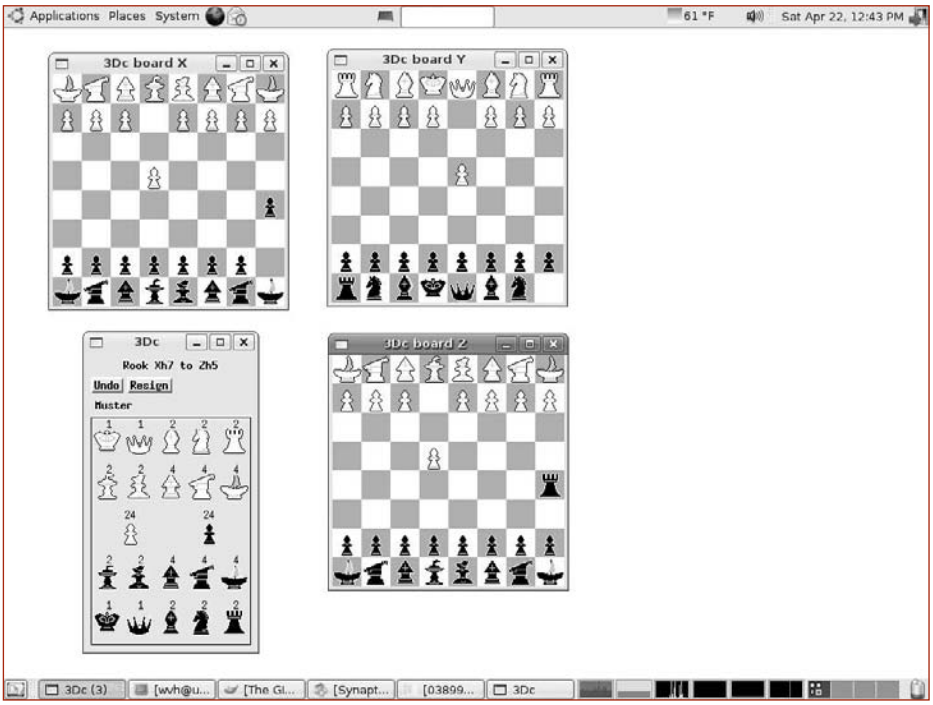
Another popular chess display application is eboard, which is available in the Ubuntu repositories and displays a nice help screen that explains how to configure `eboard` for various chess engines when you

first start the application. If you're using Kubuntu, you may also want to look at knights, which is an extremely nice KDE-oriented chess display engine.

Of course, you may have already beaten Bobby Fisher or Big Blue or simply may be bored with standard chess. For you, Ubuntu offers the 3dchess package (shown in Figure 20-6), which displays three chessboards and can automatically play against you if you start the game with the --play white|black argument, which specifies the color that you want the computer to play.

FIGURE 20-6

A 3D chess game in progress



Unless you're already familiar with 3D chess, you'll want to read the documentation in /usr/shared/doc/3dchess/3dc-rules.html so that you understand how this particular implementation of the game is played.

TIP **Fans of other versions of chess such as Japanese Shogi, will be glad to know that you can install both the `xshogi` display program and the `gnushogi` shogi-playing engine from the Ubuntu repositories.**

Go

Besides chess, one of the most historically popular strategy games in the known universe is Go, a game with a demonstrable history of several thousand years. Just like chess, Go has the advantage of being

wildly popular and free, so it matches chess in the open source space in terms of the number of free front ends for displaying games, a clever GNU Go-playing engine (gnugo), and a number of free Go servers on the Internet that you can access to play Go any time of the night or day against opponents throughout the Internet universe. A simple search for "Go" against package names and descriptions in the Synaptic Package Manager shows many packages—the ones that I'll discuss in this section are cgoban (a playing interface) and gnugo (a Go-playing engine).

Figure 20-7 shows the startup screen shown by the cgoban graphical interface for Go.

FIGURE 20-7

Initial game selection and configuration in cgoban



This screen enables you to determine if you want to play against IGS (Internet Go Server), play against NNGS (No Name Go Server), or use the Go Modem protocol to play against the gnugo program or even another human if your ISP is down. Figure 20-8 shows how to use the Go Modem dialog to specify that one of the two players is human and the other is the /usr/games/gnugo program.

Figure 20-9 shows an actual game of Go in progress after filling out additional details about a local Go game in a per-game configuration dialog that enables you to specify rule sets, board size, per-player handicaps, and komi.

FIGURE 20-8

Player specification in cgoban

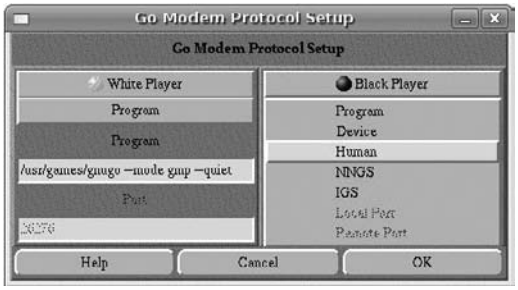
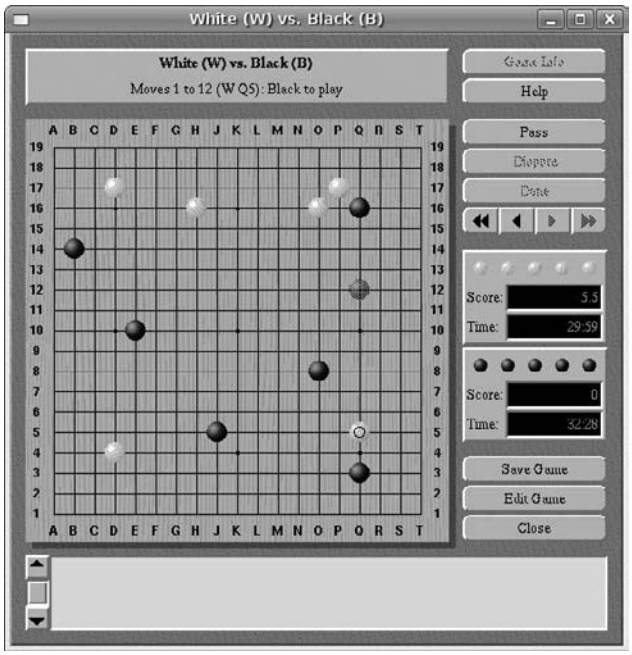


FIGURE 20-9

A Go game in progress against the gnugo engine



If you're interested in Go, the best site I've found for general information about Go is the British Go Association's (BGA) page at www.britgo.org/index.html, which provides information about applications, the history of the game, competitive Go, and a list of Internet Go servers. A great general site is the International Go federation's page at `www.nihonkiin.or.jp/igf`, which I have never found as friendly as the BGA page, but has a great set of pointers to Go sites all over the world. Sensei's Library of Go Servers at `http://senseis.xmp.net/?GoServers` also provides a fairly large list of online servers that speak various protocols.

Monopoly Clones

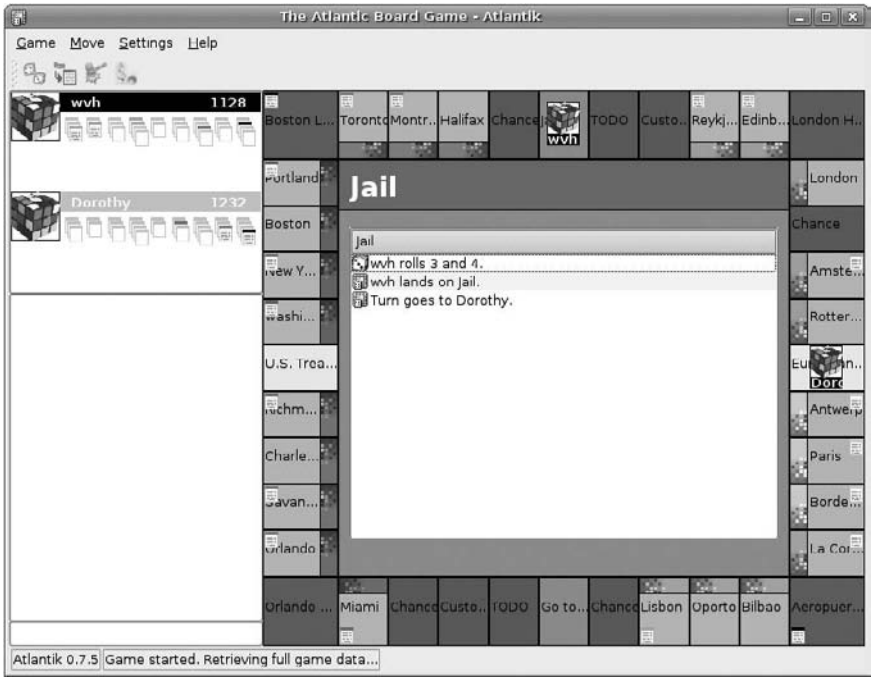
Like chess and Go, Monopoly is a board game that is incredibly popular. Unfortunately for us, Monopoly is indeed owned by Hasbro (which bought Parker Brothers, Inc., which originally developed the game—see `www.hasbro.com/monopoly`), which has not seen fit to provide a Linux port of an online version of Monopoly. However, fear not, open source fans! An excellent Monopoly-like game called Atlantik is available in the Ubuntu repository, as is a game-playing engine that enables you to play Atlantik against other players over the network. GNOME Ubuntu users will want to install the gtkatlantik package from the resulting list, while Kubuntu users will find that the Atlantik package is preinstalled on KDE 3-based Kubuntu systems (because Atlantik was originally developed for KDE). No KDE 4 version of Atlantik for Kubuntu 8.10 or newer systems was available when this book was written, but I'm sure that's an oversight that will be corrected soon. Regardless of the desktop you are using, you will also want to install the monopd package if you ever want to actually host a game using your machine as a server (which you probably will).

NOTE **Atlantik is a network-aware game that can use a server, but it does not offer automated game play against a computer opponent. This means that you will either have to contact a remote opponent from a default Atlantik server or you will have to install the `monopd` package and enable users to connect to your server to play the game.**

After you install the appropriate packages, you can start Atlantik by selecting it from the Applications ➪ Games menu on your Ubuntu system or the K Menu ➪ Games ➪ Board Games menu on your KDE 3-based Kubuntu system. Figure 20-10 shows Atlantik being played using the gtkatlantik interface.

FIGURE 20-10

A game of Atlantik in progress



Atlantik normally plays its own version of Monopoly, but you can also play the traditional Monopoly game from within Atlantik. Figure 20-11 shows a game of Monopoly in progress using the gtkatlantik interface. There are few play-related differences between this and the Atlantik game shown in Figure 20-10. The primary difference is the game board used—the rules are identical by default, although you can customize them in Atlantik's Game Configuration screen, which is discussed later.

Getting to the point where you actually see a game screen can be slightly confusing. When you first start gtlatlantik, click Enter to get past the splash screen, and a game selection dialog such as the one shown in Figure 20-12 appears.

FIGURE 20-11

A game of Monopoly in progress using Atlantik

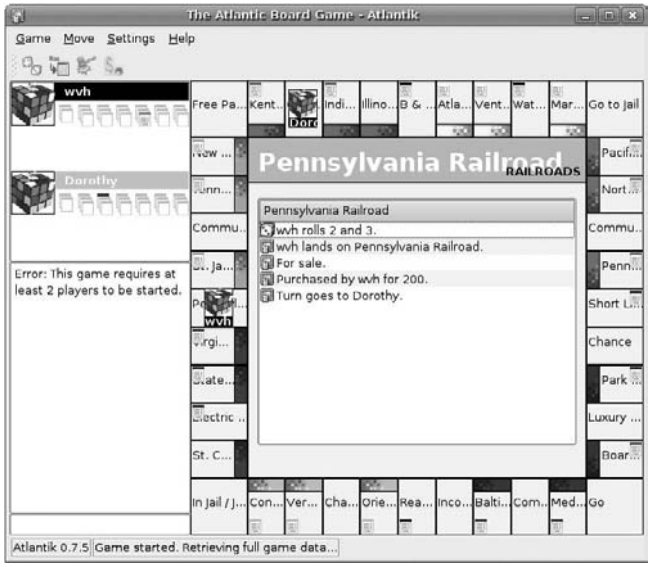
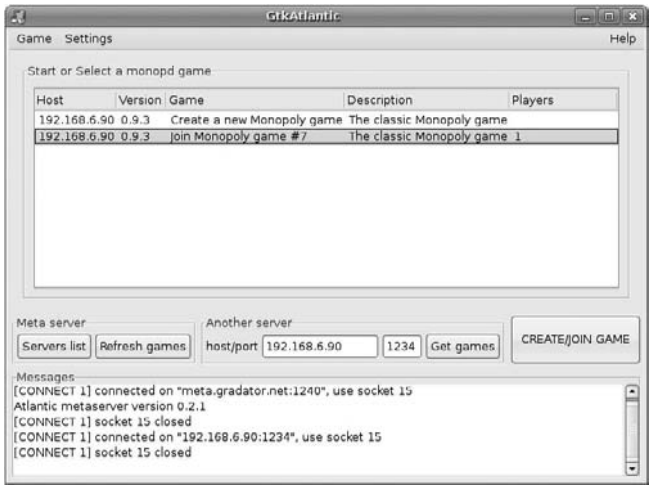


FIGURE 20-12

Creating or joining a networked game of Atlantik



Enter the hostname or IP address of the `monopd` server that you want to contact in the Another server text box. Click Get games, and a list of the games available at that server is displayed. If one is available, select that game from the list, and click CREATE/JOIN GAME, at which point the screen shown in Figure 20-10 or Figure 20-11 appears. The actual game that you see depends on the type of game(s) that the server is hosting and the one that you connected to.

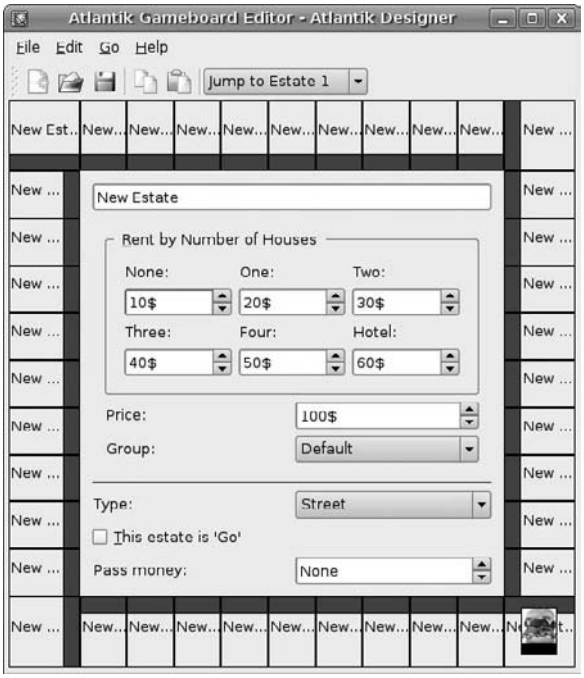
If no game is available and you are running the `monopd` server, you will need to start one and wait for someone to connect to it. This is almost the same process as joining an existing game, except that you would select one of the Create a new Monopoly game or Create a new Atlantik game entries in the game list. Your system will display a Game Configuration dialog until someone remotely connects to your game, at which point you will see the game board as the game begins.

TIP **Installing the `monopd` server package starts the `monopd` server on your system. If you cannot contact a remote `monopd` server or simply want to run your own, you can check its status using a command like `ps alxww | grep monopd`. If this doesn't show any `monopd` processes running, you can start one with a command like `sudo -u nobody /usr/sbin/monopd`.**

If you're a true Atlantik/Monopoly fan, a game board editor for Atlantik, called Atlantik Designer, is also available as a separately installable package. Figure 20-13 shows the initial Atlantik Designer screen with a blank board that you can customize, save, and use.

FIGURE 20-13

Using the Atlantik Designer to create your own game board



Using the Atlantik Designer in conjunction with Atlantik can be a lot of fun, because you can create your own game boards to reflect the streets and corporate landmarks in your own area. Who doesn't want to rule (or at least own) their neighborhood?

Atlantik is the easiest way to become a capitalist without actually crushing anyone under your boot heels, and it is great fun, just like the original Monopoly board game. Frankly, I wish that Hasbro would release an official Linux version because I believe in rewarding companies for fun games, but until then Atlantik is just all right with me.

Scrabble Clones

Scrabble is another classic board game that is extremely popular but suffers from the same commercial ownership and corporate malaise as Monopoly. Scrabble is owned by Milton Bradley, which is also a wholly owned subsidiary of Hasbro. Microsoft Windows and Macintosh versions of Scrabble are available, but until Hasbro sees fit to release a Linux version, we'll just have to content ourselves with open source clones. Unfortunately, the unofficial Linux version (PyScrabble, written in Python) that is provided in the Ubuntu repositories requires that you play against an online Scrabble server, and no suitable servers are available on the Net at the time that this book was updated. You can try to install and run your own PyScrabble server by installing and starting the `pyscrabble-server` package), but this didn't work correctly at the time that this book was written because of dependencies on a Zope database library that was not available for Python 2.5 at the time that this book was updated. If you're a Scrabble fan, you may want to install these packages and give them a try—servers may be available and the server package itself may have been fixed by the time that you read this. You have nothing to lose but a little time.

Another alternative exists for hard-core Scrabble fans. The Internet Scrabble Club (`www.isc.ro`) provides servers for its own application, WordBiz, which is a Scrabble clone that it written in Java and will therefore run on your Ubuntu or Kubuntu system. Before downloading and installing this package, you should register with their site at `www.isc.ro/en/registration.html`. After doing so, download the zip archive of their game from the URL `www.isc.ro/linux/WordBiz18linux.zip` and extract its contents using the unzip command. This will create a directory named WordBiz. Use the `cd` command to change directory to this directory and start the WordBiz application using the following command:

```
java -jar wordbiz.jar
```

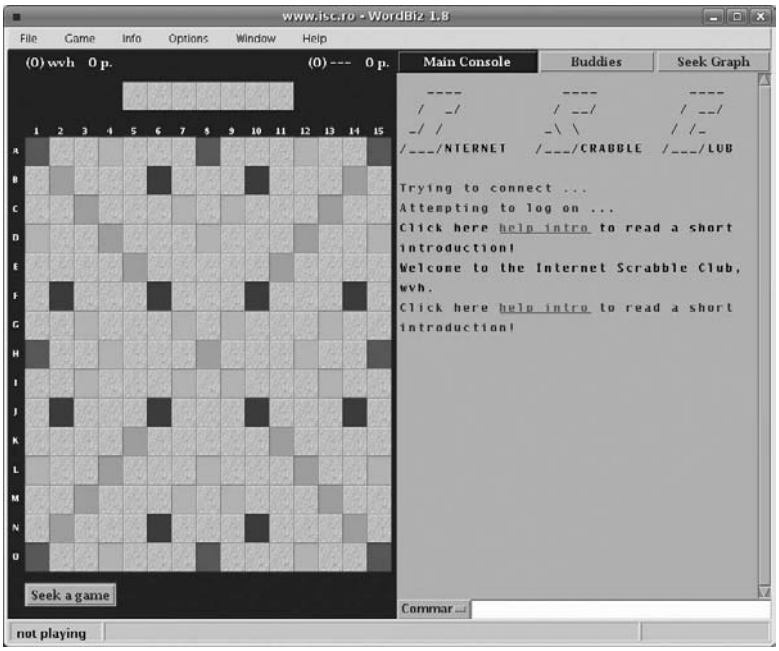
This will start the WordBiz application in its own window, shown in Figure 20-12, and will also display an Open connection dialog in which you must supply the handle (login name) and password that you registered on the Internet Scrabble Club site.

Once you have logged in, select the Game ➪ Seek Game menu command to find another person who wants to play Scrabble over the Internet, and away you go!

TIP **Other versions of Scrabble that actually provide a graphical interface are available for Linux systems, although not currently in a prepackaged form that can be easily installed on an Ubuntu system. You can get the tarballs for xscrabble, an X11-based version of Scrabble, at `ftp://ftp.ac-grenoble.fr/ge/educational_games`. A version of Scrabble called Quackle that provides its own gameplaying engine is available from `www.quackle.org`. You may also want to look at a Java-based Scrabble game called Jduplicate, which is available at `http://jduplicate.sourceforge.net`.**

FIGURE 20-14

Playing the Internet Scrabble Clubs' WordBiz game



Unique Arcade Games for Linux

The previous sections primarily discussed GNOME and KDE games that are deeply inspired by classic arcade, board, and card games. While many of these are great fun to play, they just aren't always exciting to watch or all that unique. This section highlights a few games that feature good graphics, are genuinely fun, are largely unique, and (best of all) are still provided in the Ubuntu repositories. No visit to Best Buy, Wal-Mart, or Electronic Games necessary—installing them is just a package management command away, and they are fun for the whole family.

Atomic Tanks

The Atomic Tanks arcade game is a tank fighting game that has good graphics and fun game play. You can install the Atomic Tanks game by using your favorite package management tool to install the `atanks` package. After installing this package, an entry for it is present on the Games menu on Ubuntu systems, and the Games ➪ Arcade menu on Kubuntu systems. Figure 20-15 shows a game of Atomic Tanks in progress on a Ubuntu system.

Atomic Tanks game play consists of taking turns shooting at another tank. You have a certain number of shields that you can use to protect your tank from actual hits, and missing the other tank (or hitting a shielded tank) blasts holes in the ground. As you move from level to level, you get money that you can spend on better shields and bigger weapons and great features like being able to teleport your tank to a new location.

For more information about Atomic Tanks, see its home page at `http://atanks.sourceforge.net`.

FIGURE 20-15

Playing Atomic Tanks



DOOM

Doom is a classic first-person shooter game, in which you walk the halls of a military facility haunted by daemons, aliens, and other opponents. The first rule of Doom—shoot anything that moves. In addition to the commercial versions of Doom from Id Software (some of which run on Linux—see `http://zerowing.idsoftware.com/linux/doom`), open source versions of Doom are instantly installable on your Ubuntu or Kubuntu system by installing packages such as `prboom` or (my favorite) `lxdoom`, which installs Free Dm (Free Doom). Figure 20-16 shows a game of Free Doom in progress.

If you're into fast, pulse-pounding games, Doom is, was, and will always be a classic. Playing Free Doom will capture both your cursor and your heart.

Frozen Bubble

Frozen Bubble is an open source bubble-popping game that you can play against the computer, against another local player in split-screen mode, or over a network against up to four other players. (Some of these features are only available for Linux systems, which is a pleasant change!) Your character is a penguin (Frozen Pingus) who shoots hanging balloons to form groups of balloons of the same color. Figure 20-17 shows a game of Frozen Bubble in progress.

FIGURE 20-16

Playing Free Doom on Ubuntu

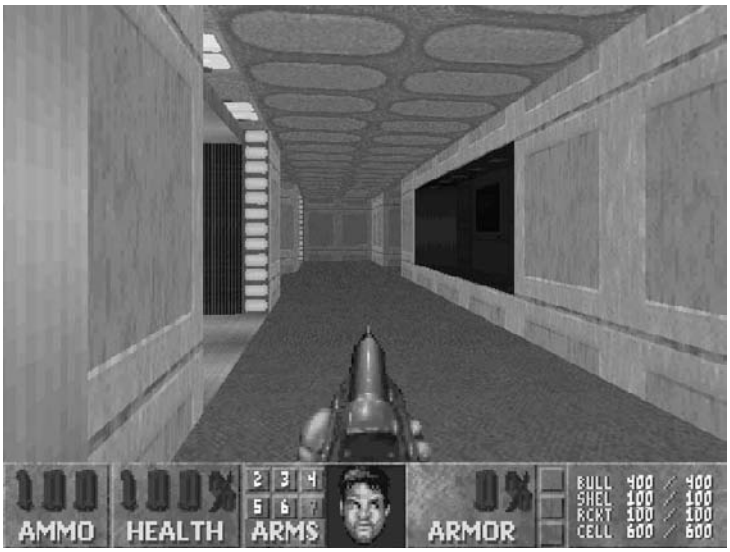
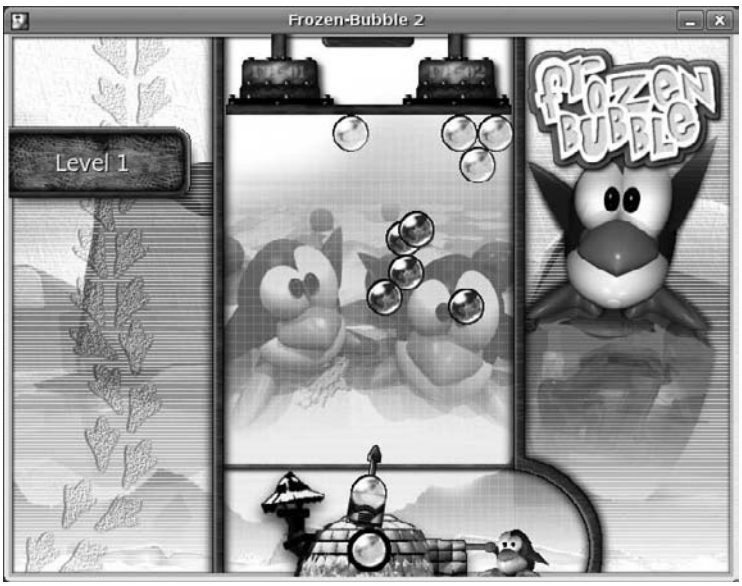


FIGURE 20-17

Playing Frozen Bubble



Frozen Bubble was inspired by 1990s arcade games such as Bubble Bobble and Puzzle Bobble—if that means nothing to you, ask your parents. The game features excellent sound, graphics, and game play, and also includes a level editor so that you can create your own game levels. For more information about Frozen Bubble, see its home page at `www.frozen-bubble.org`.

Metal Blob Solid

Metal Blob Solid, or "Blob Wars" as it is affectionately known by its aficionados, is an impressive and original Linux arcade game. You install Metal Blob Solid by installing the blobwars package using your favorite package management software. After installing this package, an entry for it appears on the Games menu on Ubuntu systems, and the Games ➪ Arcade menu on Kubuntu systems.

After starting the game, you are the character Bob, a Blob freedom fighter who is fighting to free his Blob home world from aliens who have invaded it both to assimilate the population and to find four mysterious Crystals—Fire, Time, Space, and Reality (which is, coincidentally, the title of a song by the Moody Blues). When playing the game, you primary jobs are to rescue imprisoned compatriots, shoot aliens, and pick up treasures along the way. Just to mix things up a bit, Figure 20-18 shows a game of Metal Blob Solid in progress on a Kubuntu 8.10 system.

FIGURE 20-18

Metal Blob Solid on a Kubuntu system



Metal Blob Solid is a truly fun game that could easily have sucked my pockets dry back in the heyday of arcade games. Given that it is free and available for my desktop, there is absolutely no reason to ever stop playing.

If you like Metal Blob Solid, you should investigate the other great Linux games that are available from the same folks at Parallel Realities (`www.parallelrealities.co.uk`). Project: Starfighter and Virus Killer are two of my personal favorites, but they're all great.

Tux Racer

Tux Racer is a 3D action game that features Tux, the penguin mascot of Linux, in various downhill racing scenarios, collecting herrings for extra points as he goes. You can increase Tux's speed by navigating so that you are sliding on ice, and you must avoid obstacles such as trees and rocks.

Tux Racer has an interesting history because it began life as GPL software, was marketed as a product containing closed-source enhancements, and then came back to life through multiple forks of the last versions of the original GPL code. For this reason, several different versions of Tux racer are available, including SuperTux (a 2D version, in the `supertux` package), Planet Penguin Racer (in the `planetpenguin-racer` package), and Extreme Tux racer (in the `extremetuxracer` package). Figure 20-19 shows a game of Extreme Tux Racer in progress.

FIGURE 20-19

Playing Extreme Tux Racer



Tux Racer is fast moving, attractive, and fun. The Extreme Tux racer variant is the only version that is still actively being developed and supported.

More Great Arcade Games

In addition to the default games provided with both Ubuntu and Kubuntu systems and the few extra games highlighted in this section, a tremendous number of other open source games are available in the Ubuntu repositories. This chapter could go on forever—instead, I'll just highlight a few other games

that I've had great fun playing. Some other popular games that you can install directly from the Ubuntu repositories are:

- **Armagetron:** A faithful implementation of the original Tron arcade game that puts simplified versions (like KTron) to shame. Install the armagetron package to play the local or networked versions of this game - it's great! (`www.armagetronad.net`)
- **Battle for Wesnoth**: A fun, fantasy-based strategy game that features many different story lines. Fight to regain the throne of Wesnoth, ally yourself with Dwarves or Orcs to conquer the world, or be evil yourself and use your undead army to try to conquer the world. The game includes a tutorial to get you started, after which you begin to build an army of different fighters and proceed to conquer. You install this game by installing the wesnoth package. (`www.wesnoth.org`)
- **Sauerbraten**: A great first-person shooter in the Doom tradition, but with great graphics and some interesting effects. Install this game by installing the sauerbraten package. (`http://sauerbraten.org`)
- **Torcs**: A fun 3D racing simulation that you can install by installing the OpenGL package. (`http://torcs.sourceforge.net`)
- **Wormux**: The mascots of various free software programs battle each other using dynamite, grenades, bazookas, and baseball bats. Wormux is a fun 2D game that you can install by installing the wormux package. (`www.wormux.org/wiki/en/index.php`)

The games listed in this section are just the tip of the Linux iceberg. A quick Web search for something like "Linux Games" will show hundreds more. You may never do real work again!

Summary

Both the Ubuntu and Kubuntu desktop systems provide a great selection of computer games, ranging from classic card and board games to more modern video games. Beyond the default set installed with an Ubuntu desktop system, the Ubuntu repositories provide many, many more, including some very sophisticated, graphically oriented video games and simulations. Whether you want to play games like Go or Chess, which have been around for thousands of years, or you simply want to shoot aliens, Ubuntu and Kubuntu can keep you occupied and happy until you finally have to give in to sleep deprivation or the need to go to work.

Chapter 21 discusses how to retrieve and install files on consumer electronics devices such as iPods and other music players, digital cameras, and PDAs. The types of consumer electronics devices are basic accessories for most people today, but they aren't very useful if you can't update their contents or, in the case of a digital camera, retrieve your current photos to free up space for capturing your next vacation or family get-together.

Chapter 21

Consumer Electronics and Ubuntu/Kubuntu

Linux is the embedded operating system of choice for many modern *embedded devices*. Embedded devices is the industry term for things that have computers or control systems running inside them. Twenty years ago, popular embedded devices were things like power plants, factories and assembly lines, and spacecraft. Today, the list still features those same technological chestnuts, but now also includes things like your car, your home gateway or router, your TiVo or other digital video recorder (DVR), your personal digital assistant (PDA), your MP3 or other digital audio player, and so on. Today, many consumer electronics devices contain a full-blown operating system to support their configuration and communication requirements.

Although the fact that Linux is the operating system of choice for most new embedded devices may not matter to you personally, it is very significant for you as a consumer and as a user of a desktop Linux distribution. There are a few major reasons for the popularity of Linux in the embedded world. The fact that Linux is free and, therefore, requires no payment of royalties or licensing costs is extremely important to embedded hardware vendors, who are planning on making a few million devices. First, if they had to pay some sort of royalties on the operating system they used, that cost would ordinarily be something that they would have to pass on to the consumer (you and me). Second, the flexibility of Linux makes it easy to adapt to a huge number of devices, which means that the support in Linux for the device drivers that control those devices and their peripherals continues to expand. And finally, the openness of Linux makes it easy for Linux to access files stored in almost any modern type of filesystem and communicate with just about any other operating system. This is good news if, for example, a company is building a new digital audio player wants the freedom and power of Linux but needs their customers to be able save and sync their personal audio collections to it from their Mac OS X or Microsoft Windows systems. Read and write support for popular Windows filesystems? It's in there. Support for USB connections, wireless connections, or serial ports? Free. I think that you see where I'm going here…

IN THIS CHAPTER

Recognizing consumer electronics devices

Connecting to a digital camera

Working with iPods and other MP3 players

Working with CompactFlash and SD cards

The fact that you and I are using Linux on the desktop puts that same power and flexibility in our hands. As you'll see in this chapter, you can talk to just about anything from a Linux system. Users of other operating systems require that the makers of specific consumer electronics devices or third-party software vendors provide the software that enables an operating system to talk to a specific device. Microsoft Windows users would find their iPods to be nicely designed paperweights if Apple didn't provide iTunes for Windows. PDA users who use Macs and Windows boxes would find their digital assistants to be remarkably deaf and dumb if Apple, Palm, and various third parties didn't provide synchronization software for them. Sometimes this software is free, and sometimes it costs money—but never the latter for you and me.

As you'll see in this chapter, your Ubuntu or Kubuntu Linux system can talk to, sync, and exchange files and data with almost anything. In the world of Linux, all it really takes is a motivated hacker with the latest widget from Spacely Sprockets, and sooner or later you'll be able to download and install a free sprockinator package that will make it easy for you to keep that widget in sync with your Ubuntu Linux system. It may require a bit of tweaking to do so, but you can never beat the price.

Configuring Your System for Consumer Electronics Devices

The section of Chapter 19 on "Setting CD and DVD Preferences" explained how to configure your system's behavior when a blank, audio, or data CD or DVD is inserted. CDs and DVDs may be the type of removable media that you'll use most frequently with your computer system, but consumer electronics (CE) devices are probably a close second. The following sections discuss how to configure how your system handles digital cameras, PDAs, digital audio players, and so on. (How your system interacts with removable media such as flash drives (USB sticks) and external drives is discussed in detail in Chapter 24, "Adding Hardware and Attaching Peripherals.") The next two sections focus on what your system does when you attach various types of consumer electronics devices to your Ubuntu or Kubuntu system, and how you can configure your system's behavior.

Digital Device Recognition on Ubuntu Systems

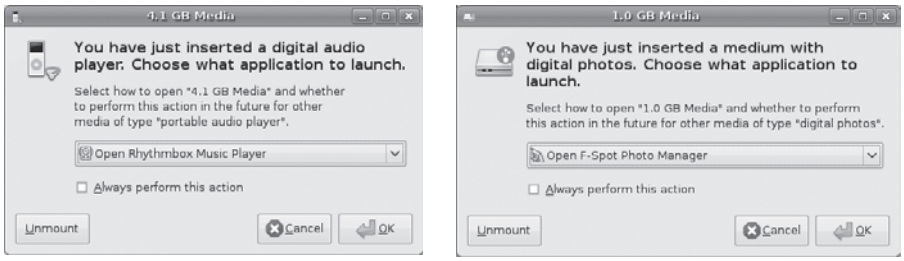
Digital devices such as music players, digital cameras, mobile Internet devices, and smart phones are a key element of today's culture. Though these devices are designed to work in isolation, they almost always require some sort of connection to a desktop or laptop to synchronize, update, or archive content. Each type of device has its own connection requirements. For digital music players, such as iPods and MP3 players, connecting to a laptop or desktop system enables you to synchronize libraries and playlists or simply to add new music. For digital cameras, you are usually downloading and archiving photos to free up space on the device. For smart phones and PDAs, you're usually synchronizing address books and archiving notes and other content.

Default Actions

When you attach a digital device such as a music player or digital camera to your Ubuntu system, it displays a dialog like the ones shown in Figure 21-1. These dialogs suggest a default action that Ubuntu believes is appropriate for that device—the dialog on the left is from a sample system to which I attached an MP3 player, while the dialog on the right was displayed after I attached a digital camera to the same system.

FIGURE 21-1

Prompt dialogs for music players and digital cameras



Prompt dialogs such as the ones shown in Figure 21-1 suggest the default action of opening an application that is associated with the new media or device. These example dialogs suggest using the Rhythmbox application to access the files on a digital audio player and suggest using the F-Spot application to access the files on a digital camera. If you're comfortable with always being prompted for what to do when a specific type of media is attached to your system and want to use the suggested application, you can simply click OK to start that application. This doesn't change GNOME's default behavior when your system detects new media of that type—the next time you insert the same device or another device of that type, the same prompt will be displayed.

If you want to perform some other action for the new media or device but want to be prompted the next time you insert the same device or another device of that type, left-click on the drop-down list of possible actions to display a list such as the one shown in Figure 21-2.

FIGURE 21-2

Possible actions for a music player

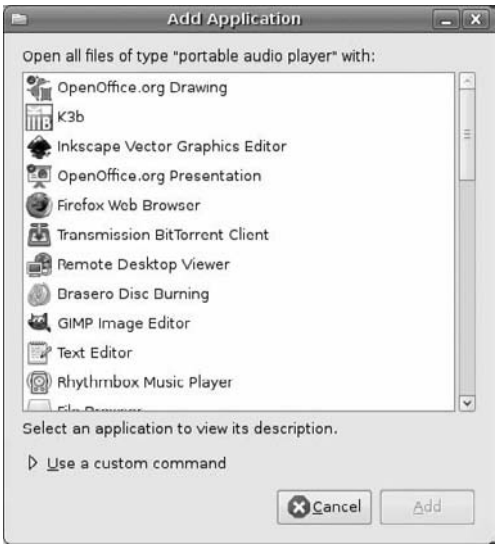


- **Do Nothing:** GNOME displays a desktop icon for the newly attached device, but does not open that device in the Nautilus file manager or any other application.
- **Open Folder:** GNOME displays a desktop icon for the newly attached device and displays the contents of that device in the Nautilus file manager.

- **Open Rhythmbox Music Player:** GNOME displays a desktop icon for the newly attached device and opens that device in the Rhythmbox Music Player application.
- **Open with other Application**: GNOME displays a desktop icon for the newly attached device and displays a dialog that enables you to select or browse for another application with which to open devices of this type, as shown in Figure 21-3.

FIGURE 21-3

Selecting another application to use with a music player



If you select Open with other Application to choose another application to associate with the current type of device, selecting an application from the dialog shown in Figure 21-3 doesn't actually open the device in the new application. Instead, it simply adds the selected application to the drop-down menu of possible options for that type of device. For example, Figure 21-4 shows the drop-down menu for digital audio devices after having selected Sound Converter as another application to use with devices of that type.

Regardless of which of the items that you select from the drop-down menu in dialogs such as those shown in Figure 21-1, you will see the same dialogs the next time that you attach the same devices. If you're comfortable with always being prompted for what to do when a specific type of media or device is attached to your system and simply want to start working with your device, you can skip ahead to the section of this chapter that discusses working with that type of device on an Ubuntu system. However, if you always want to use a specific application with a specific device or type of device, proceed to the next section for information about customizing the behavior of your system when a specific device or type of device is attached.

Customizing Digital Device Recognition

Many Linux distributions, including Ubuntu 8.04 and earlier, provided special configuration utilities that required you to uniquely configure how your system reacted when different types of consumer electronics devices were attached. For example, Ubuntu 8.04 and earlier provided the System ➪ Preferences ➪ Removable Drives and Media Preferences application. It was convenient to be able to use one application to do this, but it was still a pain because it didn't take advantage of common underlying features of these devices, such as the fact that almost all of them contain a filesystem that is stored on flash memory. You don't really need special applications to handle the devices themselves, just their content.

FIGURE 21-4

Your new choices after selecting another application to use with a music player



The latest versions of GNOME, such as the version (2.24) provided with Ubuntu 8.10, recognize this fraternity among digital devices. Configuring a specific application for use with various consumer electronics devices is now done by setting preferences in the Nautilus file manager, in exactly the same way that you configure what happens when you insert a CD or DVD.

To configure how your Ubuntu system reacts when you attach an MP3 player, iPod, digital camera, or other type of device to your system, start Nautilus by clicking on an existing filesystem icon or by selecting any entry from the GNOME desktop's Places menu. Select the Edit ➪ Preferences menu item, and select the Media tab to display the pane shown in Figure 21-5.

TIP **You can also start GNOME's Preferences application from the command line by executing the `nautilus-file-management-properties` command.**

This dialog enables you to customize what happens when different types of media are inserted in or attached to your system. As you can see from Figure 21-5, the default behavior of Nautilus is to prompt you for what to do when a DVD video is inserted in your system, when a music player is attached, when a device containing photos is attached, or when media containing installable software is attached. Figure 21-1 showed examples of these of prompt dialogs.

To permanently change what happens when a specific type of device (Music Player) or media containing certain types of files (Photos, Software) is attached to your system, left-click on the drop-down list to the right of the appropriate label in Figure 21-5 to display the list of options for that device or the types of files that it contains. Figure 21-6 shows the list of choices for music players (i.e., portable audio devices).

FIGURE 21-5

Default settings for consumer electronics devices in GNOME



FIGURE 21-6

Possible default actions for music players



This list is similar to the list that you see when clicking the drop-down menu on one of the prompt dialogs shown back in Figure 21-1, which was shown in Figure 21-2. However, the list shown in Figure 21-6 enables you to specify your system's default behavior, so it also includes an Ask what to do option. As you can see from Figure 21-6, you can specify the following default actions in response to attaching a digital audio player to your system:

- **Ask what to do:** GNOME displays a desktop icon for the newly attached device, and displays a prompt dialog like the ones shown in Figure 21-1 that enables you to specify what to do for each device of this type.
- **Do Nothing:** GNOME displays a desktop icon for the newly attached device, but it does not open that device in the Nautilus file manager or any other application.
- **Open Folder:** GNOME displays a desktop icon for the newly attached device and displays the contents of that device in the Nautilus file manager.
- **Open Rhythmbox Music Player**: GNOME displays a desktop icon for the newly attached device and opens that device in the Rhythmbox Music Player application.
- **Open with other Application:** GNOME displays a desktop icon for the newly attached device and displays a dialog that enables you to select or browse for another application with which to open devices of this type.

Selecting any of these options changes how your system will respond to devices of a given type whenever one is attached in the future.

Changing your system's default behavior when a specific type of device is attached is a good idea if you always use a specific application to work with devices of that type. For example, if you always want to use F-Spot to extract and manage photos from your digital camera, then setting F-Spot as the default action for Photos is probably a good idea. Similarly, if you always want to use another application, such as Google's Picasa, to work with your photos, you may want to select Open with other Application and select that application from the list. If the application that you want to use isn't shown on the list, click Use a custom command and enter or browse for the location of the command that you want to use.

On the other hand, you will not want to change your system's default behavior if you use different devices of specific types and want to open each in a different application. For example, I'm a big fan of the `gtkpod` application for working with iPods, but I have a couple of more generic MP3 players that I use from time to time. I, therefore, don't set `gtkpod` as my default application for music players because it has no idea how to work with generic MP3 players. Instead, I use Open with other Application from the menu and dialog shown in Figure 21-2 to add `gtkpod` to the list of applications associated with music players, and then select the appropriate application from the prompt dialog whenever I attach an iPod or other MP3 player.

Digital Device Recognition on Kubuntu Systems

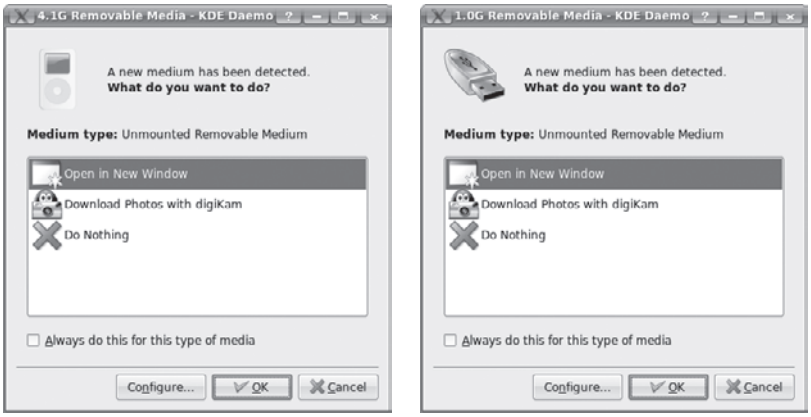
As you might expect, different versions of Kubuntu behave differently when consumer electronics devices are attached, because they are based on different versions of KDE, with different default behaviors and customization opportunities. The next two sections discuss how KDE 3-based Kubuntu systems (Kubuntu 8.04 and earlier) and KDE 4-based systems (Kubuntu 8.10 and later) behave by default when digital device such as music players and digital cameras are attached, and the extent to which you can customize that behavior. KDE 3-based systems such as Kubuntu 8.04.

Default Actions on KDE 3-based Systems

KDE 3-based systems such as Kubuntu 8.04 and earlier recognize the underlying characteristics of digital devices such as digital cameras and portable audio players by identifying them as a Removable medium that contains storage (usually flash of one type or another). Kubuntu 8.04 displays dialogs like the ones shown in Figure 21-7 when you attach a digital camera or portable audio player.

FIGURE 21-7

KDE 3 default dialogs when attaching a digital device with associated storage



The left portion of Figure 21-7 shows the dialog and associated menu that is displayed when you attach a portable audio player, while the right portion of Figure 21-7 shows the dialog and associated menu that is displayed when you attach a digital camera. While both of these as recognized as Removable storage mediums of some sort (and thus use the same dialog and menu by default), they are still recognized as being different—note the different icons in the upper-left portions of each dialog. I'll explain how to customize this menu later in this chapter in the section entitled "Customizing Digital Device Recognition on KDE 3-based Systems"—for now, let's just work with the default entries.

TIP

These dialogs contain a checkbox labeled Always do this for this type of media, which can be both very useful and very irritating. Activating this checkbox is useful if the device that you've connected is the only, or only type of, removable media device that you ever plan to attach to your system. Selecting this option, choosing a menu entry, and clicking OK will cause KDE to skip the notification dialogs when you attach a device of this type and automatically perform the action associated with the selected menu item.

On the other hand, activating this checkbox can be very irritating if you subsequently attach another Removable Medium device that actually has a different purpose. A good example of this is selecting this checkbox, selecting the Download Photos with digiKam menu item, and then clicking OK. If you subsequently attach a portable audio player, which KDE also views as a Removable Medium, KDE will try to open your music player in digiKam, which is certainly not the right tool for the job. If you set this option and then find that you regret your actions, you can simply edit the file `~/.kde/share/config/medianotifierrc` to remove the line that defines the default action for `media/removable_unmounted` or `media/removable_mounted`—or simply delete the file entirely if you simply want to start over. I refuse to confess how long it took me to figure this out the first time I made this mistake.

If you have attached a digital camera to your Kubuntu 8.04 (or earlier) system and you want to manage your photos using digiKam, select the Download photos with digiKam menu entry and click OK to start digiKam. See the section on "Using digiKam," later in this chapter, for information about importing, organizing, and editing your photos. If you don't want to use digiKam but simply want to browse your photos and copy them somewhere, you can choose the Open in New Window option, which will mount the storage medium in your camera on your system and open the filesystem on your camera's storage medium in Dolphin. You can then drill down through the DCIM directory until you find the photos that you are looking for, and do whatever you want with them.

If you have attached a digital audio player, you can choose either the Open in New Window option, which will mount the storage medium in your audio player on your system and display the filesystem on your audio player in Dolphin. This is fine on non-Apple audio players, which usually simply have a DOS-style filesystem on their storage media, in which the music is either stored in a single directory or in multiple subdirectories. The way that your music is stored and organized on Apple's iPods is a whole different ball game. See the section on "Digital Audio players and Amarok on Kubuntu," later in this chapter, for more information.

Nonstandard Camera Interfaces on Kubuntu Systems

Though most digital cameras nowadays show up as storage devices when you attach them to a Kubuntu system, some use proprietary USB interfaces to access the storage that they provide. If you attach a camera with a proprietary USB interface to a Kubuntu 8.04 or earlier system, you will see a dialog like the one in the following figure when you attach your camera.



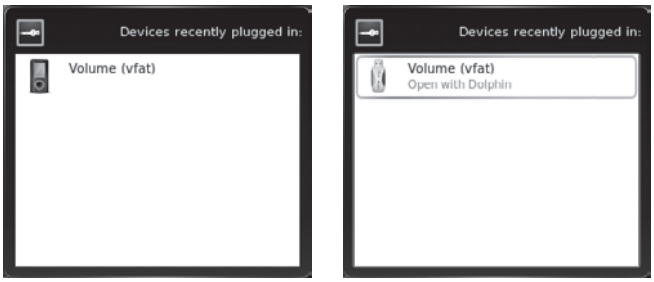
If you see this dialog, you should select the Open in New Window option to cause your Kubuntu system to mount the camera's internal storage on your system. You can then use the Cameras section of DigiKam's Settings ➪ Configure DigiKam dialog to add an entry for that camera, or simply use DigiKam's Import ➪ Browse media to browse for the device.

Default Actions on KDE 4-based Systems

KDE 4 takes a simple and straightforward approach when digital devices are attached to your system, popping up a dialog that identifies the type of the device and provides a best guess of the type of content that it contains. Figure 21-8 shows two examples of this pop-up dialog, the one on the left for a digital music device, and the one on the right for a digital camera. As you can see from the image on the right, moving the mouse over the entry for a device displays an Open in Dolphin message—to mount that device and open it in the Dolphin file manager, simply left-click on the name of the device.

FIGURE 21-8

KDE 4 notifications when attaching a digital device



As you can see from these dialogs, KDE 4 simply recognizes newly attached digital devices as storage devices. By default, it does not currently offer options for automatically starting an application (other than Dolphin) when a specific type of device is attached, but you can force it to do so by creating your own startup files, as described in the section entitled "Customizing Digital Device Recognition on KDE 4-based Systems," later in this chapter.

Customizing Digital Device Recognition on KDE 3-based Systems

The dialog at the left in Figure 21-7 shows the default menu that displays when you attach a digital audio player to a Kubuntu 8.04 or earlier system, while the dialog at the right in Figure 21-7 shows the default menu that displays when you attach a digital camera to a Kubuntu 8.04 or earlier system. These dialogs differ only in the icon that they contain—otherwise, they're the same, which is somewhat disappointing, since they seem to think that "all the world's a camera," and therefore don't provide an option for automatically starting Amarok to interact with a digital audio player. Luckily, the KDE 3 desktop environment used on Kubuntu 8.04 and earlier systems makes it easy to add new options to the dialogs that it displays when new devices are attached.

To add the menu entry to the dialogs shown in Figure 21-7 that enables you to start Amarok when an appropriate digital device is attached to your system, click the Configure button. The screen shown in Figure 21-9 appears.

To add a new menu entry, click Add. A dialog like the one shown in Figure 21-10 appears.

FIGURE 21-9

Adding or editing options for handling digital devices

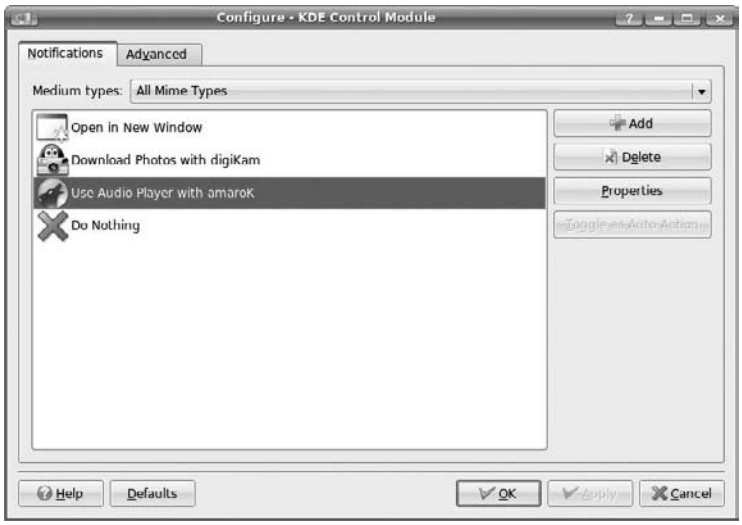


FIGURE 21-10

Defining a new menu entry for handling digital devices

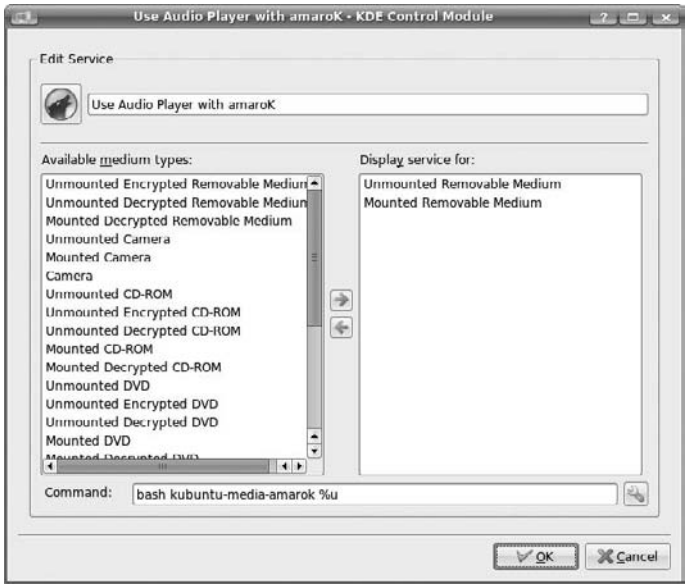


Figure 2-10 shows the values that you might want to enter when creating a menu entry for starting Amarok that will be displayed whenever a mounted or unmounted removable storage medium is attached to your system. The icon, top, and bottom text fields, and the right pane are all blank when you begin creating a new menu entry. To create an entry like this one, do the following:

1. Enter the text that you want to appear on the menu in the field at the top of this dialog. I think of this as the name for the menu entry, even though the field isn't labelled as such.
2. Click the icon field to the left of this field to display a Select Icon pop-up dialog that enables you to select an icon to display beside the new menu entry. Once you've located an appropriate icon, select it and click OK to close this dialog. The dialog shown in Figure 2-10 is updated to show the new icon to the left of the name of the entry that you are creating.
3. In the left pane, select each medium that the new menu entry should be triggered by, and click the right-facing arrow to add that medium to the list shown in the right pane. For digital audio devices, I typically specify that this menu item should be present on the dialog that opens when either mounted or unmounted removable media are detected.
4. Enter the command that you want to execute when the new menu entry is selected. The command shown in Figure 21-10, `kubuntu-media-amarok`, is simply a shell script that starts Amarok in a mode that begins playing or accessing music from the current media device (specified by the `%u` parameter).
5. Click OK to save the new menu entry. The next time that a medium of the specified type is attached to your system, the new menu will also be displayed in the dialog that opens, as shown in Figure 21-11.

FIGURE 21-11

The new menu entry in the New Medium Detected dialog



TIP **If you want to set up custom device notification dialogs for new media detection in KDE 3-based systems, you don't have to wait until you actually attach a new device to do so. You can manually display the device recognition dialog for new media at any time by executing the `kdesu kcmshell media` command and entering your password.**

Customizing Digital Device Recognition on KDE 4-based Systems

As mentioned elsewhere in this book (especially in Chapter 8), KDE 4 is a work in progress, which means that some things happen faster than others. One thing that has not happened very quickly is making it easy for users to customize what happens when you attach a digital device such as a USB stick, digital camera, or digital audio player to your system. At the moment, the only option in the Device Notifier for any of these is "Open with Dolphin" (as shown in Figure 21-8), which is not the appropriate action in many cases and is downright irritating if you're not a fan of the Dolphin file manager.

The mechanism for enabling applications to interact with newly attached digital devices on KDE 4-based systems such as Kubuntu 8.10 and later is that such applications will install a file in the directory `/usr/share/kde4/apps/solid/actions` with a name of the form *`something`*`.desktop`, where *something* would usually be the name of the application (but doesn't actually have to be). The following is an example of a file that would offer to start the Dragon media player when a DVD disk was inserted into your system's DVD drive:

```
[Desktop Entry]
X-KDE-Solid-Predicate=[ StorageVolume.ignored == false \
      AND OpticalDisc.availableContent == 'Data|VideoDvd' ]
Type=Service
Actions=open;

[Desktop Action open]
Name=Open with Dragon Player
Exec=dragon --play-dvd
Icon=dragonplayer
```

NOTE **The `X-KDE-Solid-Predicate` examples in this section should be on a single line and should not include a backslash—the backslash is only present for formatting purposes in the examples in this book. Pages are only so wide, after all!**

Similarly, a file that would offer to open DigiKam when anything containing a filesystem was attached would be the following:

```
[Desktop Entry]
X-KDE-Solid-Predicate=[ StorageVolume.ignored == false \
      AND StorageVolume.usage == 'FileSystem' ]
Type=Service
Actions=open;

[Desktop Action open]
Name=Open with DigiKam
Exec=digikam --detect-camera
Icon=digikam
```

Finally, an example that would provide an option to open Amarok whenever an audio CD was inserted is:

```
[Desktop Entry]
X-KDE-Solid-Predicate=[ StorageVolume.ignored == false \
```

```
      AND OpticalDisc.availableContent == 'Audio'  ]
Type=Service
Actions=open;

[Desktop Action open]
Name=Play CD with Amarok
Exec=amarok --cdplay /dev/cdrom
Icon=amarok
```

TIP **This same Amarok entry is customized for CDs and won't handle digital audio players. After all, it's just an example. You should be able to create your own Amarok entry for digital audio players using the examples in this section. Hint: they typically show up as a device containing a filesystem.**

Unfortunately, files such as these are not installed by default in the version of KDE 4 that came with the original release of Kubuntu 8.10. Eventually, you shouldn't have to create your own `.desktop` files—applications will begin to install them automatically. Until every application installs its own `.desktop` file, you can create your own (or copy the examples from this section) in the `/usr/share/kde4/apps/solid/actions` directory. For example, after adding the three files in this section to this directory and inserting a DVD, the Device Notifier pop-up provides more interesting options, as shown in Figure 21-12.

FIGURE 21-12

The Device Notifier pop-up with additional options



Because multiple options are available, clicking the entry for the DVD that I've just inserted displays a dialog like the one shown in Figure 21-13. Look familiar?

I'm not suggesting that the device notification files that this chapter provides are optimal examples—they're just ones that I hacked together to make life more pleasant on my KDE 4 systems. (The DragonPlayer startup file is excerpted from an example in Harald Sitter's blog at `http://amarok.kde.org/blog/categories/17-apachelogger`.) Until applications actually install these types of files, feel free to use these examples or roll your own. When creating your own notification files, you can get information about the possible values in the X-KDE-Solid-Predicate test from the KDE documentation at `http://api.kde.org/4.0-api/kdelibs-apidocs/solid/html/classes.html`.

FIGURE 21-13

A customized new device dialog triggered by the Device Notifier



Working with Your Digital Camera

By this point in this book, you may have noticed that Ubuntu/GNOME and Kubuntu/KDE have completely different applications for doing the same sorts of things, which shouldn't be surprising given the fundamental differences in the desktop and development environments that they provide. As Steve Martin once noted when discussing the English language, "The French have a different word for everything!" Choice is good, and Linux is all about choice (and the freedom to make that choice, of course).

The next two sections discuss the primary applications that Ubuntu and Kubuntu provide for working with your digital camera, namely F-Spot and DigiKam. These applications were introduced in the "Photo Editing Overview" section of Chapter 18 (which also lists some alternate digital camera-related applications that are available for Ubuntu and Kubuntu), but the following sections focus on F-Spot and DigiKam in detail.

Working with Digital Cameras on Ubuntu Systems

Prior to Ubuntu Hardy (8.04), Ubuntu systems used a shell script named `gnome-volume-manager-gthumb` to use the gthumb image viewer to import, organize, and display your photographs. In Ubuntu Hardy and later, this application and the associated script are still available, but the default application for importing, organizing, and displaying digital photographs is now F-Spot, an excellent application that is unfortunately written using the proprietary and odious .NET framework that is supported on Linux thanks to the MONO project (which is apparently not named after the disease, as I have always thought). If you want information about using `gthumb`, look for a PDF copy of the first edition of this book on thepiratebay.org—this section focuses on using F-Spot. Political and nerd considerations aside, it's really a great application.

TIP **If you're upgrading an older Ubuntu system or have simply stored your existing digital photos in another directory, you can import them into F-Spot by using the Photo menu's Import command.**

When you attach your digital camera to a USB port on your Ubuntu system and turn on the camera, your Ubuntu system displays the dialog shown in Figure 21-14, identifying the fact that a device containing digital photos has been detected and asking how you want to proceed.

FIGURE 21-14

Your Ubuntu system detecting a digital camera



To use F-Spot, the recommended application, simply click OK. If you always want to automatically use F-Spot whenever any device containing digital photos is detected, you can select the Always perform this action checkbox before clicking OK. A dialog like the one shown in Figure 21-15 opens.

The dialog shown in Figure 21-15 provides two options that you may want to consider before importing the photographs from your camera:

- **Tag:** Add a customized tag to the set of photos that you importing, which will make it easier to locate this set of photos in the future. Activating this checkbox enables you to select an existing tag to apply from a drop-down list, or to create your own. The predefined tags are Favorites, Hidden, People, Places, and Events. You can also select the Create New Tag option at the bottom of the list to display a dialog in which you can define your own tag, which I often do. This lets me use tags like "Mom's Birthday," "Luke's Graduation," or any other event-specific tag that will make it easy for me to find a related set of photos in the future. You can define top-level tags like the predefined ones, or define you own tags within one of the predefined tags to help organize your photos in the way that best suits you.
- **Detect Duplicates:** Analyze the photos that you are importing and only import those that are new. This option is active by default, and can be very handy if you previously imported photos and did not remove them from your camera after importing them.

After making any changes to these settings that you want to use, click Copy to import your photos or click Cancel to close the dialog without importing your digital photographs. If you click Copy, an F-Spot dialog like the one shown in Figure 21-16 displays once your photos have been successfully imported. By default, photos are imported into subdirectories of the Photos directory in your home directory, based on the dates on which the photos were taken. For example, photos taken on November 25, 2008 would be stored in the directory `~/Photos/2008/11/25`. You can change this location using the Edit menu's Preferences dialog, which is discussed later in this section.

FIGURE 21-15

The F-Spot application preparing to import your photographs



TIP **Many digital cameras automatically power down after a short period of time with no activity. This is great for saving battery life, but may cause import failures if the camera powers down while the F-Spot application is waiting for input. If you are importing photos and F-Spot creates files that do not contain your photos, try repeating the import process.**

FIGURE 21-16

Imported photos in F-Spot



As shown in Figure 21-16, the primary F-Spot window consists of the following areas:

- **Left sidebar:** Lists the available tags if no photo is selected, and also provides details about any photo once you select one
- **Right pane:** Initially displays thumbnails of your recently imported photos
- **Find Bar:** Located above the right pane, this bar enables you to display the sets of photos based on certain criteria. For example, if you selected or created a tag when importing a set of photos, you can select that tag in the left pane and drag it to the Find bar to display only photos which have that tag. The Find menu provides other selections criteria. For example, if you want to display photos with which no tags are associated, you can select the Find menu's Untagged Photos command. You can also use the Find menu's By Rating, By Date, or By Import Roll submenus to display photos by those criteria. (Import Roll essentially identifies a set of photos that you imported from your camera at the same time, using a single import session.)
- **Toolbar:** Provides easy access to common F-Spot commands
- **Menus:** Provide easy access to all F-Spot commands

TIP **F-Spot doesn't always display its initial dialog correctly. If you start F-Spot and don't see the left sidebar, you should see a slender bar with three horizontal lines in its center directly to the right of the left window border. Select this bar and drag it to the right to display the left sidebar.**

After verifying that an import session worked correctly, you can purge the photos from your camera (or other media) by selecting the Edit menu's Delete from Drive command.

The Find bar is a very useful selection device. You'll find it useful to tag selected photos (and create your own tags) so that you can locate conceptually related sets of photos regardless of when you imported them. To locate photos by tag, simply select that tag in the left sidebar and drag it to the Find bar. (If you select the Untagged Photos menu command, you must deselect it before you can find photos by tag.) To undo a tag selection, you can either right-click on that tag in the Find bar and select Remove from search from the pop-up menu that opens, or select that tag in the Find bar and drag it out of the Find bar. If only one tag was specified, either of these actions also closes the Find bar. You can toggle whether the Find bar is displayed by pressing Ctrl+F at any time. Closing the Find bar removes any find criteria that you've currently specified from the Find bar. To refresh the select of photos that are currently displayed in the right pane, click the Refresh search icon (the circular arrow). To close the Find bar, click the Close box at the right of the Find bar or press Ctrl+F.

By default, F-Spot imports your photographs into the Photos subdirectory in your home directory. You can change the location to which your photos are imported by selecting the Edit menu's Preferences command to display the dialog shown in Figure 21-17 and modifying the Import Settings sections Folder option.

F-Spot is a great application with many features, all of which are discussed in detail in its online document, which you can display by selecting the Help menu's Contents menu option or pressing F1 while working in F-Spot. Some of my favorite F-Spot capabilities are the following:

- **Editing Photos:** Selecting a photo and clicking the Toolbar's Edit button displays a list of edit options in the left sidebar. These provide common editing tasks like copping, red-eye reduction, desaturation (decreasing the amount of a certain color in a photo), sepia toning, straightening (rotating the image to improve its alignment in the frame), Soft Focus, Auto Color (automatic color correction), and Adjust Colors. The Desaturate, Sepia Tone, and Auto Color edit options are applied immediately once they are selected. Selecting any of the other editing options displays associated controls for that type of edit in the left pane.

TIP **F-Spot does not have an Undo function. Instead, it saves multiple versions of your images. If you don't like a specific effect, you can switch to a previous version by selecting it from the Photo menu's Versions submenu. By default, the versions of a photo do not have specific names, but you can create a named version of a selected photo by selecting the Photo menu's Create New Version command, and entering a name and clicking OK in the pop-up dialog that appears.**

- **Exporting to Online Photo Services:** Selecting some number of photos and selecting an option from the Photo menu's Export to submenu enables you to export photos from F-Spot to online photo sites like SmugMug, PicasaWeb, Flickr, 23hq, and Zoomr. A dialog like the one shown in Figure 21-18 displays, explaining that F-Spot must be authorized in order to export photos to the specified online service (in this case, Flickr).

FIGURE 21-17

F-Spot's Preferences dialog



- After specifying the access that you want other users on the remote site to have to your account, click Authorize to proceed, which will start your default Web browser so that you can log in to the specified service. Once you do this, your browser will display a screen like the one shown in Figure 21-19, which enables you to authorize F-Spot to access your account and upload your photos.
- Once the upload completes, click Done to close the summary dialog and return to F-Spot.

F-Spot is a great application for importing, managing, and exporting photos. Its support for tagging makes organizing and locating photos extremely easy, and its support for exporting photos to web galleries, other folders, CDs, and the online services discussed in the previous section is incredibly convenient.

FIGURE 21-18

Exporting photos to Flickr using F-Spot

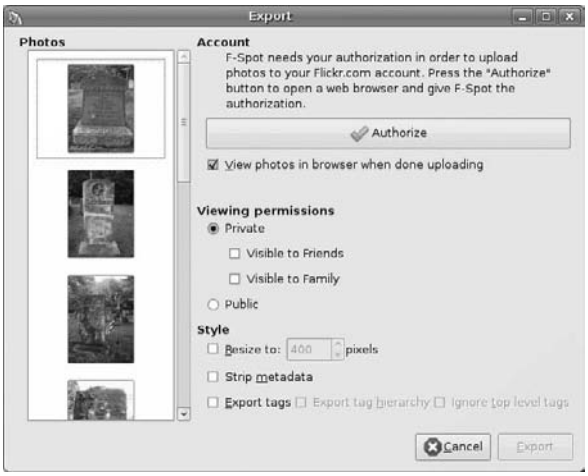


FIGURE 21-19

Authorizing F-Spot's access to a Flickr account



Working with Digital Cameras on Kubuntu Systems

DigiKam is the most popular application for working with digital photographs on a Kubuntu system, though other applications are available, most notably KPhotoAlbum (`http://kphotoalbum.org`). DigiKam makes it easy to import photos from a digital camera or other storage device, and provides a number of convenient plug-ins for common photo editing and touchup tasks.

"Open with DigiKam" is one of the default options on a KDE 3-based system when you attach a standard digital camera to a USB port on your system. (You can also add this option to a KDE 4-based system by following the instructions in the section entitled "Customizing Device Recognition of a KDE 4 System." earlier in this chapter.) You can manually start DigiKam by selecting the K Menu ➪ Graphics ➪ DigiKam - photo management menu command on a KDE 3-based Kubuntu system, or by selecting the K menu ➪ Applications ➪ Graphics ➪ Photo Management - DigiKam menu command on a KDE 4-based Kubuntu system.

If you started DigiKam manually, you can use the commands on the Camera menu to locate and browse the contents of a camera or storage device that contains your digital photos. If you start DigiKam from the Device Notifier menu after attaching a digital camera or other storage device containing photos to your system, DigiKam automatically scans your device for you. When new photos are found on a camera or other storage device, DigiKam displays the dialog shown in Figure 21-20.

FIGURE 21-20

Importing images in DigiKam



To import the images from your camera or storage device, select one of the options from the drop-down Download menu, which is available at the bottom of the dialog shown in Figure 21-20. Because I am somewhat paranoid, I typically select Download All to download the images without removing them from my camera. Selecting this option displays a dialog like the one shown in Figure 21-21, which enables you to select an existing album or create a new one into which to import your photos.

FIGURE 21-21

Selecting or creating a photo album



To import your photos into an existing album, select the name of that album. You cannot import photos directly into the top-level My Albums folder, so you'll have to create at least one new Album (the Vacations album shown in Figure 24-21 is a custom album that I created). Click OK to begin the import process. A progress bar in the dialog shown in Figure 21-20 displays as the photos are imported into the selected Album. Click Close once the import process completes to display the main DigiKam application dialog, as shown in Figure 21-22.

FIGURE 21-22

DigiKam's main screen



TIP **Once you're satisfied that your photos have been successfully imported into DigiKam, you can use the commands on the Camera menu to clean up your camera or other storage media. Select the appropriate Browse media command from the Camera menu to reexamine the device from which you imported your images. This redisplays the dialog shown in Figure 21-20. You can then select the Delete menu's Delete All command to delete all of the images on your media. Alternately, you can delete only certain images by first selecting them and then selecting the Delete menu's Delete Selected command.**

As shown in Figure 21-22, the main DigiKam screen provides some introductory information about using DigiKam, and consists of the following areas:

- **Left Sidebar and Sidebar tabs:** The sidebar enables you to view selected albums in different ways and use the sidebar tabs to search for photos by the date on which you downloaded them, by tags, or to search your photo albums. By default, the Albums tab is selected, which enables you to select a specific photo album for viewing, editing, or exporting.
- **Main Image window:** Displays the contents of a selected album or photo. You can double-click on any image to display a larger version of that image in the Image window, and can redisplay the album view by clicking anywhere in the image window.
- **Right Sidebar:** Provides icons that enables you to obtain information about your images. This sidebar provides tools for obtaining general photo properties, photo metadata, color usage, viewing and setting captions and tags, and creating tag filters for use when searching for and displaying matching images.
- **Toolbar:** Provides icons for common actions in DigiKam.
- **Menu bar:** Provides access to all DigiKam actions and dialogs.

Figure 21-23 shows DigiKam's main screen showing the contents of a specific album after that album was selected in the left sidebar.

To access any of the commands in the right sidebar, simply click on their icon. Figure 21-24 shows the Caption/Tags tool, which makes it easy to create tags and apply them to selected photos. To close this tool, click on its icon in the right sidebar a second time.

To create a tag in the Captions/Tags tool, first select the photo that you want to tag, and then select the Captions/Tags icon in the right sidebar. Next, right-click in the My Tags pane at the level where you want to create the new tag, and select New Tag from the context-sensitive menu that appears. Only the My Tags entry is available in this tool when you first use DigiKam, but you can easily create hierarchical tags that make it easy to select and view specific photos or classes of photos. For example, Figure 21-24 shows four custom tags: Dorothy, Pets, Cats, and Dogs. Cats and Dogs are subtags within Pets, so (after applying these tags to different photos) I could use the Tags tab in the left sidebar to see all photos tagged as Pets, or view specific photos tagged as Dogs or Cats.

DigiKam provides Light Table and Image Editing tools that enable you to edit and touch up your photos manually. To display the Light Table tool, select the Tools menu's Light Table menu command or press F6. A Light Table dialog appears, to which you can drag and drop images from the main DigiKam window. Figure 21-25 shows the Light table dialog after dragging and dropping a single image to it.

FIGURE 21-23

A selected photo album in DigiKam's main screen



FIGURE 21-24

DigiKam's Captions/Tags tool

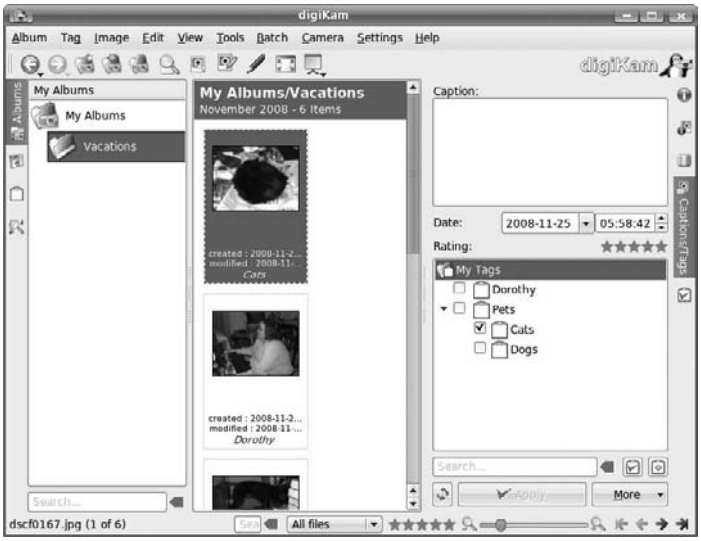


FIGURE 21-25

DigiKam's Light Table window



After displaying an image in the Light Table window, you can then right-click on an image and select Edit to open that image in DigiKam's Image Editor. You can also edit an image directly from DigiKam's main window by right-clicking on an image and selecting Edit Image from the pop-up menu that displays. Figure 21-26 shows an image in DigiKam's Image Editor.

FIGURE 21-26

DigiKam's Image Editor window



The Image Editor provides a huge selection of tools that enable you to modify color, enhance the image (including removing red eye and other favorite photo problems), transform the image by rotation, cropping, and so on, decorating an image by adding borders, applying filters for specific effects, and so on. While not quite The GIMP, DigiKam's Image Editor is a truly impressive tool. The Image Editor provides an undo command so that you can experiment with confidence, only saving your updated photo when you're happy with your changes.

If you want to edit a set of images rather than editing them one at a time, DigiKam's Batch menu provides access to a tremendous number of image conversion and modification tools that are provided by the ImageMagick package. You must separately install this package on your Kubuntu system in order to perform these effects. Once you have installed this package, you first select the images that you want to modify and then select the appropriate menu for the class of changes that you want. This displays a dialog like the one shown in Figure 21-25, which enables you to select a specific effect from the Effect menu at the top and then apply it to the selected images by clicking Start.

DigiKam is an impressive tool for importing, managing, and editing photos. If its image editor doesn't do everything you want, you can always manually edit your photos using a tool such as Krita or The GIMP—but you will rarely have to do so.

Working with iPods and Other Digital Audio Players

Digital audio players of various sorts have largely replaced other portable music players such as cassette decks and CD players. Digital audio players don't skip, are extremely small, require little power, and are hard to damage because they have few moving parts. On the other hand, when those moving parts are a disk drive and the drive goes bad, that damage is usually fatal—which makes it all the more important to have backups of your digital audio collection on your desktop machine, where it is safer (and where you hopefully have more disk space).

You can't possibly walk down the street today without seeing people who are listening to music using a digital audio player. Zillions of companies make them, the best known of which is Apple Computer could conceivably be a footnote today were it not for Mac OS X and the runaway success of its iPod digital audio player and attendant iTunes applications and music store. The iPod is a great example of what happens when people who understand industrial design meet people who understand software and usability. In many households, the iPod has replaced the automobile as the device most likely to be buffed, polished, and accessorized.

If you use a digital audio player other than an iPod, the free, open applications that are available on Ubuntu and Kubuntu systems provide complete support for downloading and managing the music on your device. However, if you're an iPod user (as I am), there is one big caveat. The iTunes application, which is the default software from Apple Computer that is used to manage music on an iPod, isn't available for Linux. The applications discussed in this section are an excellent replacement for iTunes as far as managing music on an iPod goes—but they don't enable you to buy music through iTunes. This is not a problem for me because I don't buy music there. I purchase and download MP3s from online sources that aren't tailored to a specific device, or I purchase CDs and rip them. Once you have MP3 files, you can play them on your desktop or put them on any digital audio player (including an iPod) using the applications discussed in the next few sections.

TIP **If you absolutely must use iTunes, you can actually install and run it on a Linux system using WINE, which is a software package that implements the Microsoft Windows API for Linux systems. I actually use this package when I have to run Microsoft Word, which it does a great job of. A version of WINE is available in the Ubuntu repositories, but I use a commercial version of WINE from CodeWeavers (`www.codeweavers.com`), who are the primary sponsors of the WINE project. This works great for the applications that I've tried it with, but I've never tried it with iTunes. You can also install and use iTunes in a Windows virtual machine, which are easy to set up and are discussed in Bonus Chapter 7. After all, why waste perfectly good hardware on Windows?**

The next few sections explain how to work with different types of digital audio players on your Ubuntu and Kubuntu systems. The first two sections discuss using the standard audio playback and management applications used on both Ubuntu (Rhythmbox) and Kubuntu (Amarok) systems with iPods and other types of digital audio players. Subsequent sections discuss `gtkpod`, a GTK application for working with iPods that I've used for years, and highlight some specialized applications that are tailored for other types of digital audio players, and that are available in the Ubuntu repositories.

ID3 Tag Editing on Ubuntu and Kubuntu

As discussed in Chapter 19, most of the popular audio file formats such as MP3 and OGG support ID3 tags to provide information about file contents such as the artist, album title, song title, date recorded, and so on. All digital audio players can extract and display this information from supported audio file formats, and most use it to enable you to sort, search, and organize the files on a digital audio player. Before transferring audio files to your audio player, you should verify that they have valid ID3 tags, or else it may be hard to figure out exactly who performs `Track03.mp3`. You can usually verify the ID3 tag information by simply playing your audio files in a graphical audio application such as Rhythmbox, Amarok, or xmms, because these applications automatically display ID3 information as they play each audio file.

Most CD-ripping applications will automatically fill out ID3 tag information during the ripping process, but you may occasionally encounter a CD that is not in the Internet CD databases, and for which you will therefore have to manually supply ID3 information. The Ubuntu repositories provide several applications for filling out or updating ID3 and ID3v2 tags. I have used `id3`, `id3v2`, and `id3tool`, all of which are command-line utilities and, therefore, work fine on either Ubuntu or Kubuntu systems. Others include mp3info, a terminal-oriented interactive application, and kid3, a graphical, KDE-based ID3 tag editor.

Digital Audio Players and Rhythmbox on Ubuntu

Attaching an iPod or other digital audio player to your Ubuntu system with a USB or FireWire cable displays a dialog like the one shown at left in Figure 21-1 near the beginning of this chapter. The Open Rhythmbox Music Player option is selected by default. Click OK to proceed.

When the Rhythmbox Music Player opens (you may have to click its icon in the Notification Area in the top GNOME panel to actually display it on your screen), its left pane will contain a highlighted icon for your digital audio player in the Devices section, and the right pane will contain a list of all of the audio tracks that it contains. You can click the arrow to the right of the menu item in the left pane to show all of the playlists that are currently present on that device, as shown in Figure 21-27, which lists the playlists on an iPod, but which also shows a Devices entry for a generic 4.1 GB audio player.

Playing Music

Rhythmbox will begin playing music immediately when some audio players are connected, while you must double-click a playlist or an audio track on others for them to begin playing. If you select a playlist, the tracks that it contains are listed in the right pane, as shown in Figure 21-28. If the Cover art plug-in is active in Rhythmbox and Rhythmbox can identify the album that contains the track that is playing, the cover art for that album will be displayed at the bottom of the left pane. You can enable or disable this plug-in by selecting the Edit menu's Plugins command and checking or unchecking its entry in the Configure Plugins dialog.

FIGURE 21-27

Playlists on an iPod in Rhythmbox

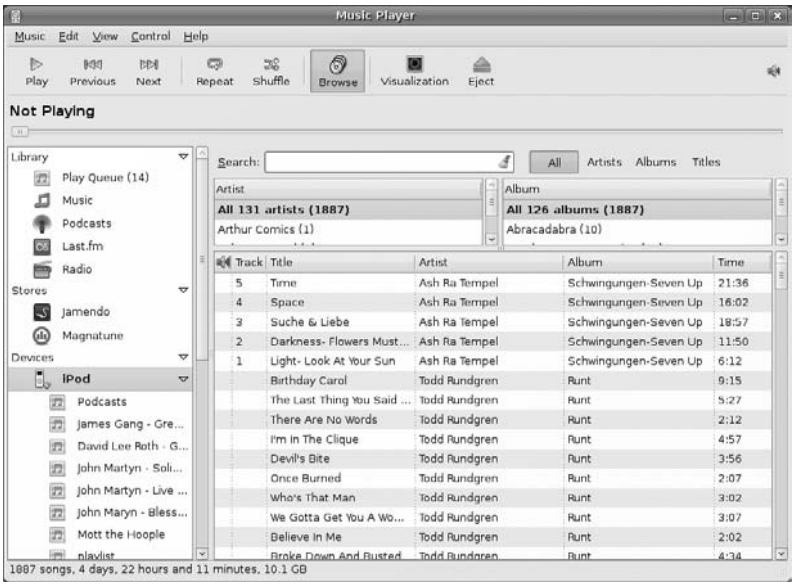
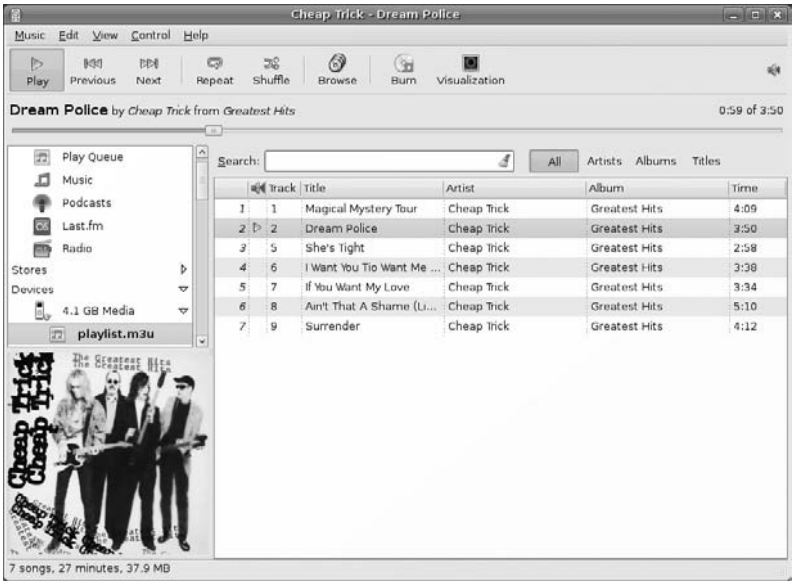


FIGURE 21-28

Playing a selected playlist



You can also play any specific song from your device by clicking the name of the device in the left pane to display all of the songs that it contains, and double-clicking a specific track to play it.

TIP **When copying or removing songs or creating and editing playlists as described in the next few sections, you can use the standard GNOME selection techniques to select multiple items at one time. To select noncontiguous entries from a list, you can hold down the Ctrl key while you select each individual entry. To select contiguous entries from a list, you can select the first entry in the range that you want to select, press the Shift key, and select the entry at the other end of the range to select those entries and all of the entries between them.**

Copying Music to an Audio Player

To copy files to a digital audio player, select the music library or device in the left pane that contains the tracks that you want to copy. The list of the songs in that library or on that device displays in the right pane. Select the songs that you want to copy, drag them to the device to which you want to copy them, and release the left mouse button.

After you have dragged songs to your audio device, a progress bar that shows the status of the copy displays in the lower-right portion of the Rhythmbox window.

Copying Music from an Audio Player

Rhythmbox also makes it easy to copy files from your portable audio player to an existing library. This can be handy if you store your music collection on several different systems or want to listen to your audio files on a laptop or machine at work without always having to carry around and use a specific portable audio player.

To copy files from a portable audio player to a Rhythmbox music library on a desktop or laptop system, click the entry for that device in the Devices section of the left pane in Rhythmbox and select the tracks, artist, or album listing whose tracks you want to copy in the right pane. Left-click and hold the mouse button down, drag your selection(s) to the Music entry for your Library in the left pane, and release the left mouse button. A status message is displayed at the bottom of the Rhythmbox window as the files are copied to your library.

Removing Music from an Audio Player

To delete files from a digital audio player, select that device in the left pane to display a list of the songs in that library or on that device displays in the right pane. Select the songs that you want to delete, and right-click to display a context-sensitive menu. Select the Move to Trash command, which will remove them from the listing. The songs will be purged when you actually disconnect your audio player, as described later in this section.

Creating, Renaming, and Editing Playlists

Playlists provide a convenient mechanism to group together and organize the songs that originally appeared together on a CD, or to simply create a list of songs that you like to hear together and in a certain order. In order to create a playlist on a digital audio player, the songs that you want to add to it must already be present on that device.

To create a playlist, right-click on the name of the digital audio player on which you want to create it, and select New playlist from the pop-up menu that displays. How you name a new playlist works slightly differently depending on the type of audio device that you are using:

- On most MP3 players, an empty playlist is displayed at the bottom of the list of available playlists. Enter the name that you want the new playlist to have—but enter it carefully because you cannot

edit playlist names once you have created them on many generic digital audio players. To rename a playlist on most audio players, you have to create a new one with the correct name, copy the entries from the old playlist to the new one, and then delete the old playlist.

- On iPods, an empty playlist named New playlist is simply added to the end of list of available playlists on your iPod. This is fine, because you can edit the names of playlists on an iPod. To rename a playlist on an iPod, right-click its name in the left pane and select Rename from the pop-up menu that displays. You can then edit the name of the playlist in the left pane and press Enter/Return on your keyboard or click elsewhere to save the new name.

Once you have created a new playlist, click the name of the device on which you created it to see the songs that it contains. You can drag songs to the playlist one by one, or use the techniques described in the tip earlier in this section to select and drag multiple songs at the same time.

Disconnecting Your Audio Player

You should not simply disconnect a digital audio player from your computer system. There should actually be no problem in doing so if you have not made any changes to its contents, but it's a good idea to get in the habit of always using the Rhythmbox Eject command before disconnecting a digital audio player. Don't worry, this doesn't actually shoot your iPod across the room—the Eject keyword is shared with devices that you actually do eject, such as CDs and DVDs. Using this command guarantees that all pending changes have been written to your digital audio device, including removing any tracks that you selected for deletion.

To eject a digital audio player in Rhythmbox, right-click on the name of the device in the left pane and select Eject from the context-sensitive menu that appears. If you have made recent changes to the device, you will see a notice in the lower-right corner of your GNOME desktop stating that data must be written to the device before you can remove it. When all data has been written to the device, another dialog opens that states that it is now safe to disconnect the device.

Digital Audio Players and Amarok on Kubuntu

Amarok is an excellent music player and organizer for your online music collection. Using Amarok with CDs and online music was discussed in Chapter 19; this section focuses on using Amarok with digital audio players.

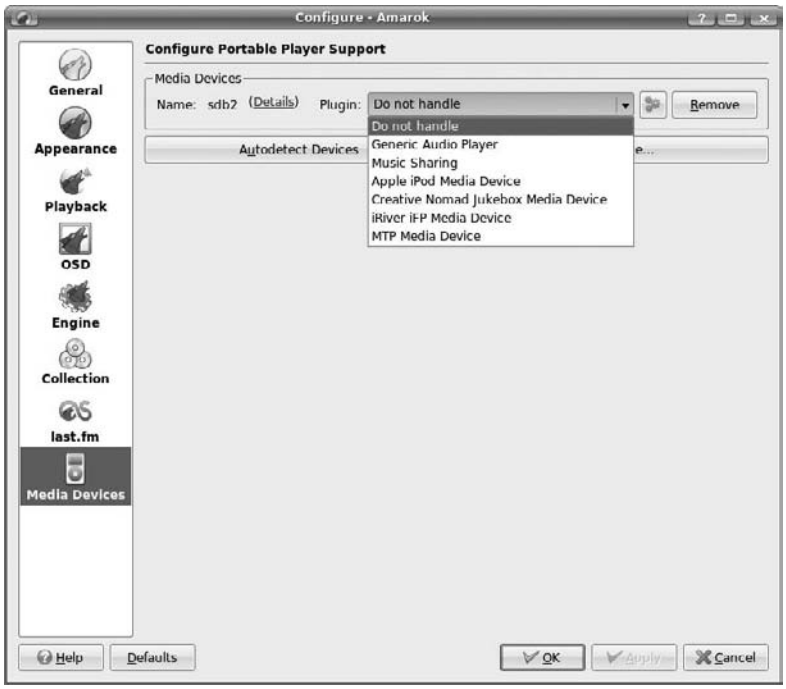
Configuring Portable Audio Devices in Amarok

Though opening a portable audio device in Amarok is an option on the default Device Notifier dialog, Amarok requires some configuration before you can easily work with portable audio devices. This is primarily so that Amarok can work correctly with audio devices such as the iPod, which store data in special formats and locations. Regardless of whether you start Amarok from the K Menu ➪ Multimedia ➪ Amarok - Audio Player menu on a KDE 3-based system, from the K Menu ➪ Applications Multimedia ➪ Amarok - Audio Player menu on a KDE 4-based system, or from the Device Notifier dialogs that were discussed earlier in this chapter, you must configure Amarok to know how to handle each media device that you want to use with it.

To configure a media device in Amarok, select the Settings ➪ Configure Amarok menu command. Select the Media Devices icon in the sidebar of the configuration dialog to display the Configure Portable Player Support dialog. The Media Devices list at the top of this dialog lists all of the portable audio devices that are currently detected. To specify the plug-in to use to deal with a specific audio player, click the plug-in drop-down list and select from the available plug-ins, as shown in Figure 21-29.

FIGURE 21-29

Configuring portable audio devices in Amarok



Once you have selected the appropriate plug-in, click OK to close the configuration dialog and return to Amarok. Click the Devices tab at the left edge of the screen. The plug-in and audio device that you selected in the Configure Amarok dialog should be listed at the top of the Media Device Browser pane. Click Connect in the upper-left portion of this pane to connect to the digital audio player. Figure 21-30 shows a connection to an iPod that has been detected as the device `/dev/sdb2`.

TIP **If you have problems connecting to a portable audio player on a KDE 3-based (Kubuntu 8.04 or earlier) system, right-click on its icon in the upper-left corner of the screen and select Mount from the context-sensitive menu that displays. You can then click Connect again in Amarok, and the connection should be successfully established. On a KDE 4-based (Kubuntu 8.10 or later) system, you may need to mount your audio player before you can use it in Amarok. The easiest way to do this at the time that this book was written was to select the entry in the Device Notifier and open it in the Dolphin File Manager, which mounts the device.**

As shown in Figure 21-30, Amarok's left pane shows all of the playlists and music found on your portable audio player.

Playing Music

To begin playing music in Amarok, expand the playlists section and select a specific playlist, or expand the entry for one of the artists whose music is stored on the device and select an Album or track to begin playing it. Figure 21-31 shows a selected playlist being played in Amarok, while Figure 21-32 shows a selected album being played.

FIGURE 21-30

An iPod in Amarok

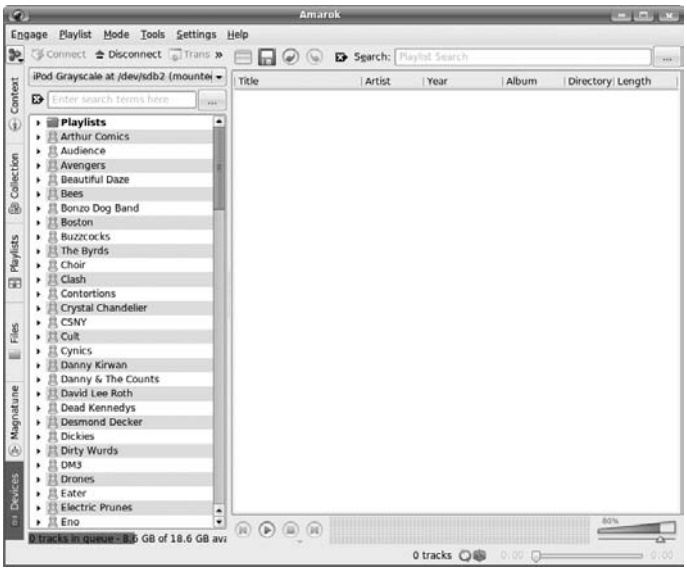


FIGURE 21-31

Playing a selected playlist

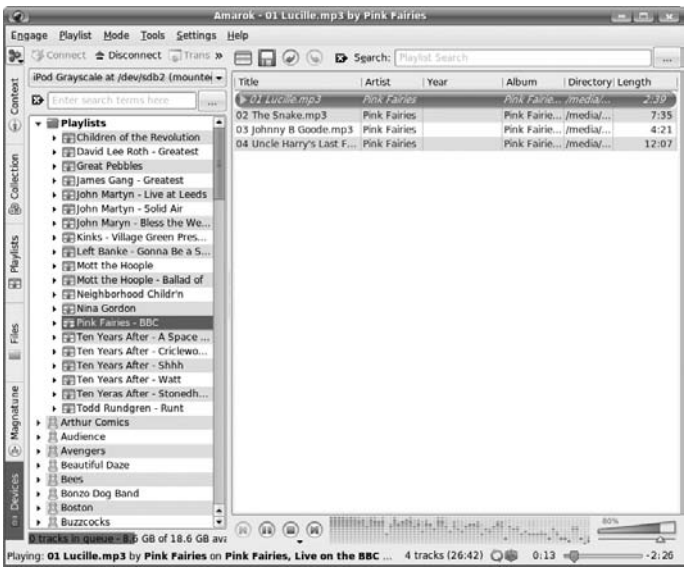
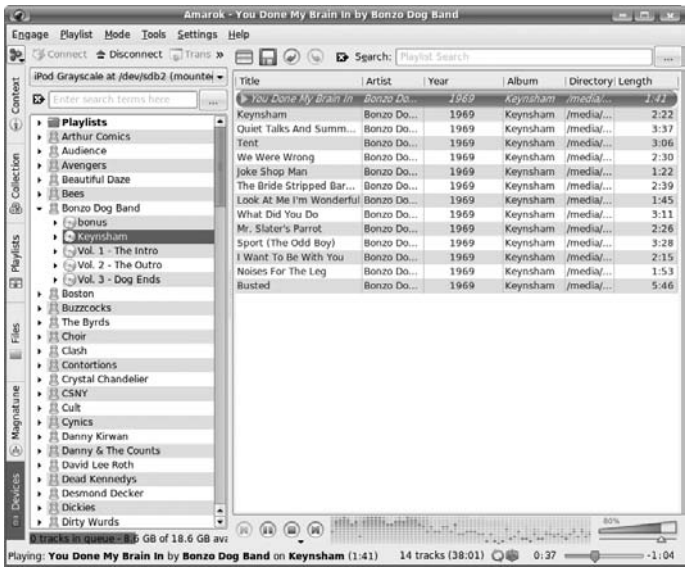


FIGURE 21-32

Playing a selected album



TIP **When copying or removing songs or creating and editing playlists as described in the next few sections, you can use the standard KDE selection techniques to select multiple items at one time. To select noncontiguous entries from a list, you can hold down the Ctrl key while you select each individual entry. To select contiguous entries from a list, you can select the first entry in the range that you want to select, press the Shift key, and select the entry at the other end of the range to select those entries and all of the entries between them.**

Copying Music to an Audio Player

To copy files to a digital audio player, select the Collection tab to display your collection in the left pane. Select the songs that you want to copy in the left pane and copy them to your audio player in one of two ways:

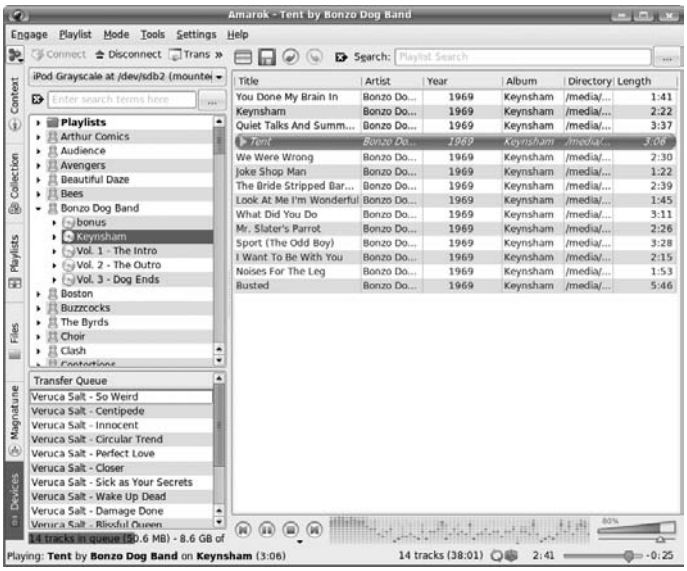
- Drag them to the Devices tab, and release the left mouse button
- Right-click on an album or artist whose files you want to copy to the portable audio player, and select Transfer to Media Device from the context-sensitive menu that opens

A Transfer Queue pane displays below the entry for your audio device, as shown in Figure 21-33.

Amarok uses a transfer queue to simplify copying audio tracks from a collection to a device, which can be a somewhat slow process and requires database updates on devices such as an iPod. Using a queue enables you to select multiple tracks or albums for copying to your portable audio player and coordinate all of the associated updates.

FIGURE 21-33

Queuing tracks for transfer to a portable audio player



After you have queued all of the albums or songs that you want to copy to your audio player, click Transfer in the top-right portion of the left pane to begin the transfer process. A progress bar that shows the status of the copy is displayed at the bottom of the left pane, and tracks are removed from the Transfer Queue as soon as they are transferred. The Transfer Queue pane closes when all files have been transferred.

Copying Music from an Audio Player

Amarok also makes it easy to copy files from your portable audio player to an existing collection. This can be handy if you store your music collection on several different systems, or want to listen to your audio files on a laptop or machine at work without always having to carry around and use a specific portable audio player.

To copy files from a portable audio player to a collection on a desktop or laptop system, click the Devices tab to ensure that the files on your portable audio player are available, and select the artist or album listing whose tracks you want to copy. Right-click the selected entry and select Copy Files to Collection from the context-sensitive menu that appears. A dialog like the one shown in Figure 21-34 opens.

After reviewing the options on this dialog (and modifying any that you want to change), click OK to begin transferring the selected files from your digital audio player to your collection. A progress bar is displayed at the bottom of the Amarok screen as the transfer proceeds. The progress bar disappears when the transfer is complete.

Removing Music from an Audio Player

To delete files from a digital audio player, select the Devices tab to display the list of the songs and artists on that device in the left pane. Select the artists or songs that you want to delete, and right-click to

display a context-sensitive menu. Select the Delete N Tracks from the Device command, where N is the number of selected tracks and device is the name of the device that contains them. Figure 21-35 shows this menu entry for an iPod.

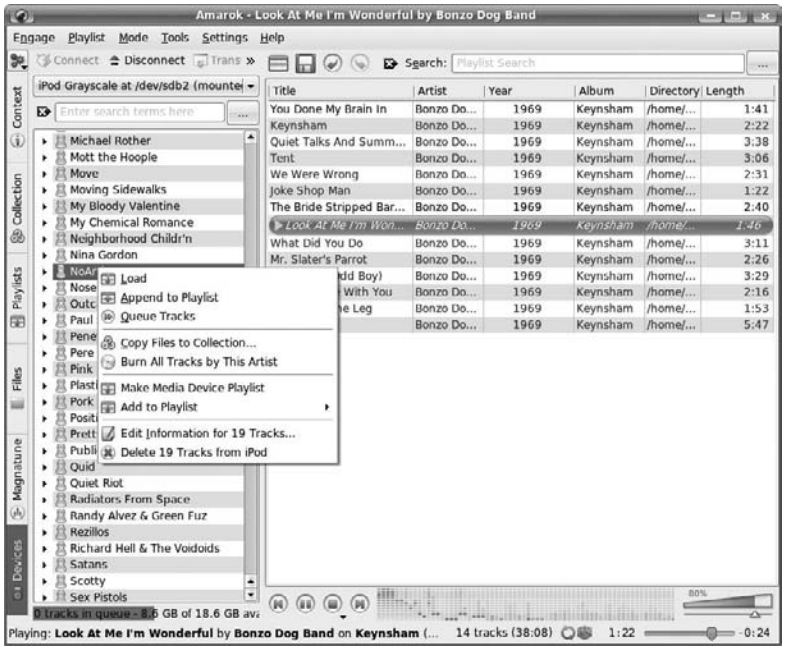
FIGURE 21-34

Copying music to a collection



FIGURE 21-35

The context-sensitive menu for deleting music



After selecting this menu entry, Amarok displays a dialog warning you that the selected tracks will be irreversibly deleted from the audio player. Click Delete to proceed, or click Cancel to abort the delete operation. If you select Delete, the selected track(s) will be immediately deleted from your audio player.

Creating, Renaming, and Editing Playlists

Playlists provide a convenient mechanism to group together and organize the songs that originally appeared together on a CD, or to simply create a list of songs that you like to hear together and in a certain order. In order to create a playlist on a digital audio player, the songs that you want to add to it must already be present on that device.

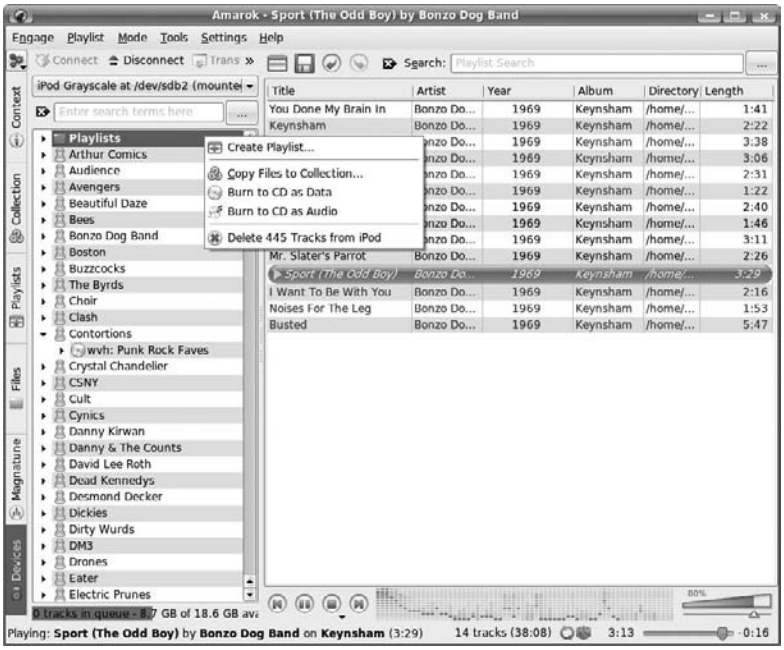
NOTE **Not all digital audio players use or support playlists—some simply play the tracks from a selected album or artist in order. If you do not see a Playlists entry in the left pane after displaying the contents of your portable audio player, and no Make Media Device Playlist entry is shown on the context-sensitive menu that is displayed when you right-click on an artist or album in the left pane, you're out of luck as far as playlists are concerned.**

You can create playlists in Amarok in two different ways:

- To create an empty playlist so that you can subsequently add songs to it, right-click the Playlists entry at the top of the left pane and select Create Playlist from the pop-up menu that appears, as shown in Figure 21-36. An empty playlist with the name New Playlist is created, with the name in edit mode so that you can change it to something more evocative of the playlist that you are creating.

FIGURE 21-36

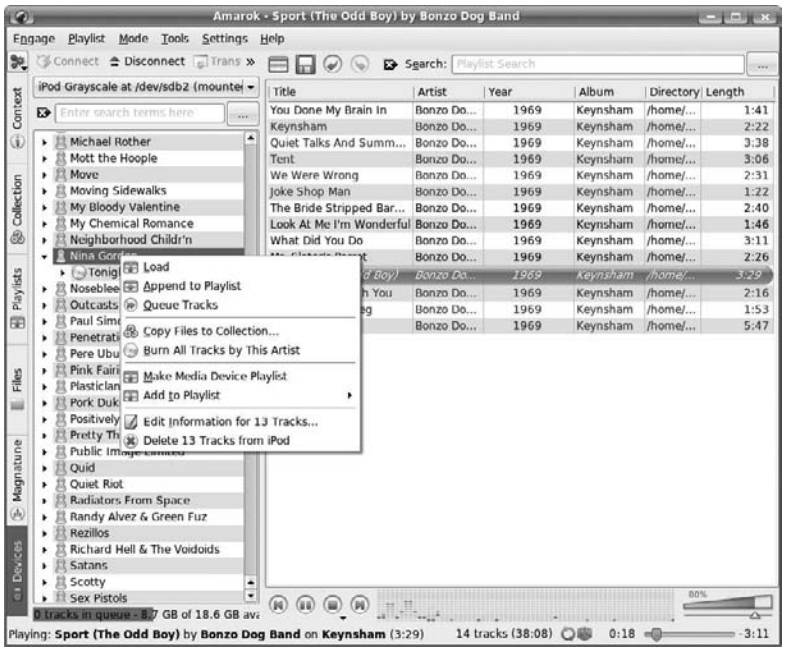
Creating an empty playlist in Amarok



- To create a playlist for a specific Album, click the Device tab to ensure that the audio device is shown on the left pane, and right-click on the artist or album for which you want to create a playlist. Select Make Media Device Playlist from the pop-up menu that appears, as shown in Figure 21-37.

FIGURE 21-37

Creating a playlist for a given selection in Amarok



If you select an artist, the new playlist is created with the name of that artist. If you select multiple audio tracks or a specific album, the new playlist is created with the generic name New Playlist. You can then right-click on that playlist and select the Rename menu command to enable you to edit the name of the playlist.

Once you have created a new playlist, click the name of the device on which you created it to see the songs that it contains. You can drag songs to the playlist one by one, or use the techniques described in the tip earlier in this section to select and drag multiple songs at the same time.

Disconnecting Your Audio Player

You should not simply disconnect a digital audio player from your computer system. There should actually be no problem in doing so if you have not made any changes to its contents, but it's a good idea to get in the habit of always using Amarok's Disconnect command before disconnecting a digital audio player. Using this command guarantees that all pending changes have been written to your digital audio device, including removing any tracks that you selected for deletion.

To disconnect a digital audio player in Amarok, select the Devices tab to ensure that the contents of your audio device is shown in the left pane. Click the Disconnect command at the top of the left pane, and wait until the listing for the device is no longer displayed in the left pane. On a KDE 3-based system, you may also need to right-click on the icon for the device in the upper-left corner of the screen and select Safely Remove from the context-sensitive menu that appears. On a KDE 4-based system at the time that this book was written, the best solution for guaranteeing that everything has been correctly updated on your digital audio device is use a file manager such as Dolphin or Nautilus to see if it is still mounted on your system, and unmount it by right-clicking on its entry and selecting the Safely Remove command from the context-sensitive menu that opens.

Working with an iPod Using gtkpod

You may be wondering why I'm even bothering to discuss a separate application for working with audio files on the iPod, when Rhythmbox and Amarok already do a great job. This is simply because that wasn't always the case—iPod support is a fairly recent addition to both Rhythmbox and Amarok. Since `gtkpod` has been around for years and still offers some capabilities that Rhythmbox and Amarok do not, I still think it's a great application for working with an iPod, especially if you don't like Rhythmbox or Amarok. So feel free to skip this section—or read on if you're interested in reading about an excellent alternative to working with your iPod using the "standard" Ubuntu and Kubuntu audio applications.

Since Rhythmbox and Amarok are installed by default and work with iPods, the `gtkpod` application is not installed by default on Ubuntu or Kubuntu systems. In order to follow along in this section, you'll have to use your favorite package management software (as discussed in Chapter 23) to install the `gtkpod` package. If you're installing `gtkpod` on a Kubuntu system, be aware that it will also install all of the GNOME/GTK libraries that it requires to run.

Attaching Your iPod

Attaching an iPod to your Ubuntu, Kubuntu 8.04 and earlier, or Kubuntu 8.10 and later system with a USB or FireWire cable displays a dialog like the ones shown at left in Figure 21-1, 21-7, and 21-8, respectively. If you are just experimenting with `gtkpod`, you can select Do Nothing from the menu on this dialog on an Ubuntu or Kubuntu 8.04 system. Your iPod will show up as a removable storage device and a desktop icon for it will automatically be created. On a Kubuntu 8.04 system, you will need to right-click on this icon and select the Mount command from the context-sensitive menu that is displayed. On a Kubuntu 8.10 or later system, you should click the entry in the Device Notifier for the iPod so that it opens in Dolphin, which mounts the device for you. If you decide that you like `gtkpod`, you can always create a custom menu entry for it in the device recognition dialogs, as described in the sections on customizing digital device recognition on Ubuntu and Kubuntu systems, earlier in this chapter.

CAUTION **When you attach an iPod to your Linux system, the iPod's screen will display an international No symbol and a Do not disconnect! message. Believe it. Follow the instructions in the section called "Safely Disconnecting Your iPod," later in this section, to safely remove your iPod without risking damage to its disk and the files that it contains.**

Using gtkpod

Installing `gtkpod` on a Ubuntu system creates the Applications ➪ Sound and Video ➪ gtkpod menu item, the K Menu ➪ Multimedia ➪ gtkpod menu entry on Kubuntu 8.04 and earlier systems, and the K Menu ➪ Applications ➪ Multimedia ➪ gtkpod menu entry on Kubuntu 8.10 and later systems. Selecting this menu item starts `gtkpod`. If your iPod was mounted and can be accessed successfully, it should be listed in the pane at the left of the `gtkpod` window. To see the actual contents of your iPod and begin working with it, select its name from this pane. The other `gtkpod` windows will update to show the

contents of your iPod, as shown in Figure 21-38. As you can see from this figure, the screen shots in this section were all taken on a standard Ubuntu system. With the exception of the window decorations, `gtkpod` looks the same on Ubuntu and Kubuntu systems.

FIGURE 21-38

The contents of an iPod in gtkpod



TIP **If you do not see your iPod listed in the left `gtkpod` pane, select the Edit menu's Edit Preferences command. In the dialog that appears, click Set mountpoint or edit repository options and verify that the entry in the iPod Mount Point field is the directory where your iPod is actually mounted and change it if necessary. If you change this field, you will have to exit from and restart gtkpod to have it check the new mount point. You may also have to click the Read icon on the `gtkpod` toolbar to force the gtkpod application to read the database information from an iPod that is mounted in a new location.**

The next few sections discuss how to do some common tasks with `gtkpod`. You can, of course, do much more—after all, there are entire books on iTunes—but the suggestions in the following sections will get you started and are the most common tasks that I do in `gtkpod`.

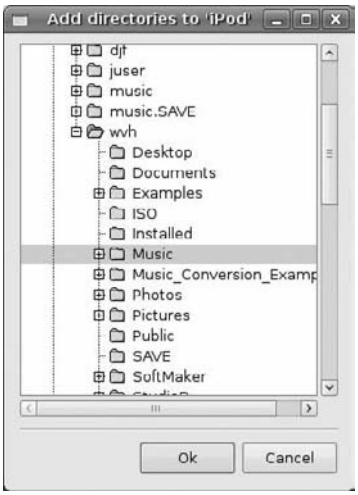
Adding Music to Your iPod

You can add MP3 files to your iPod in two different ways, either as individual files or by adding complete directories of MP3 files. To add individual files to your iPod, click the Add Files button in the gtkpod menu bar to display a standard browse dialog in which you can navigate to the directory containing the file(s) that you want to add, select them, and click Open. The file(s) that you selected will be added to your iPod, and the Artist and Album information will be added to the iPod (if it is present in ID3 tags in the audio file).

More commonly, you will want to add all of the audio files associated with a single album to your iPod at one time after you've ripped your latest CD acquisition. To do this, click the Add Folder button in the gtkpod menu bar to display the dialog, shown in Figure 21-39.

FIGURE 21-39

Adding directories of MP3 files to your iPod



Navigate to the directory that you want to add, and click Open to add that entire directory to your iPod. The main gtkpod window updates to reflect the addition of the new directory (entire album).

Even though the main gtkpod dialog shows the files and directories that you've added to your iPod, only the files have been copied there. The databases that your iPod uses to track file, song, and artist information aren't actually updated until you resynchronize the gtkpod database with the iPod's database by clicking Sync, as described in "Writing Changes to Your iPod" later in this section.

Creating and Editing Playlists on Your iPod

Creating playlists that contain your favorite music du jour is part of the fun of having and using an iPod. At a minimum, you'll probably want to create a playlist for each album that you add to your iPod, so that you can recreate the experience of hearing the tracks in the order that the artist intended.

The `gtkpod` application makes it easy to create many different types of playlists. There are also various ways of creating playlists, but you can access all of them from the Edit ➪ Create Playlists menu, shown in Figure 21-40.

For example, to create a playlist that contains all of the tracks associated with a certain album, navigate to that album in `gtkpod` and select all of the tracks that it contains. Once you've done that, select Edit ➪ Create Playlists ➪ Containing Selected Tracks. This displays a dialog that enables you to specify a meaningful name for the new playlist (the default value is the catchy name "New Playlist," which isn't all that memorable). After entering a new name, click OK to save your playlist under its new name.

As you can see from Figure 21-40, `gtkpod` supports a wide variety of playlists. Once you've created a new playlist, its name is displayed in the leftmost pane in the main gtkpod dialog. (You may need to click on the entry for your iPod in order to see the playlists that it contains.) You can then select any playlist to show the sequence of tracks that it contains, as shown in Figure 21-41, which displays my newly created playlist for the ancient Steve Miller Band album, Number 5.

FIGURE 21-40

Different types of playlists in gtkpod

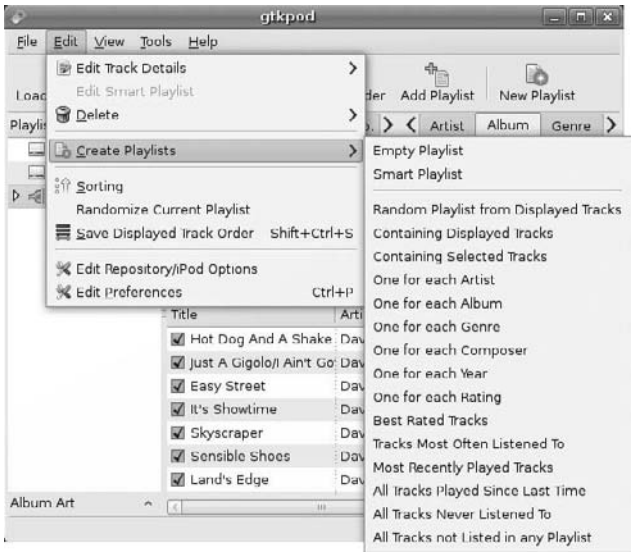
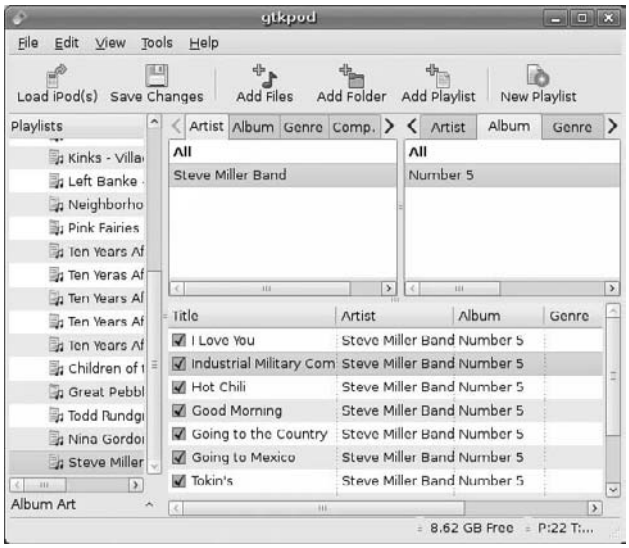


FIGURE 21-41

Examining a specific playlist



You can rename a playlist at any time by clicking its name in the left pane and entering a new name, or you can delete it and start over by selecting Edit ➪ Delete ➪ Selected playlist or right-clicking on the playlist and selecting Delete but keep tracks from the context-sensitive menu that displays. More often, you'll simply want to modify an existing playlist. To remove tracks from a selected playlist or otherwise manipulate the files that one contains, right-click on any entry in the playlist to display a context-sensitive menu. To remove a track from the playlist without removing the associated audio file from your iPod, select the Delete From Playlist command. To remove the selected track from both the playlist and the iPod, select the Delete From iPod command.

You can manipulate existing playlists several other ways. For example, you can re-sort the playlist in different ways by clicking on any of the headings shown in the track-listing pane at the bottom of the `gtkpod` application. You can add other tracks to an existing playlist by navigating to those tracks and dragging them from the track-listing pane to the name of the playlist in the left pane. The `gtkpod` application provides a tremendous amount of control over playlists, making it easy for you to create playlists that enable you to hear what you want, when you want, in the order that you want.

Writing Changes to Your iPod

Once you've finished adding music to your iPod, creating and editing playlists, and so on, you will certainly want to save all of those changes to your iPod. Although files and directories of audio files, new playlists, and other new items are added to your iPod at the time that you create them, the database that your iPod uses to track these new and modified items isn't updated until you explicitly synchronize the information in `gtkpod` with those databases on your iPod. You should always synchronize `gtkpod` with your iPod after making any changes to ensure that those changes are not lost or saved in some incomplete, intermediate state. After all, updating playlists or adding new music to your iPod doesn't do you much good if you can't access the new information!

To synchronize this information, click the Save Changes button in the gtkpod toolbar. A dialog displays as the information is synchronized. Once this dialog closes, it is safe to disconnect your iPod from your Ubuntu system, as described in the next section.

Safely Disconnecting Your iPod

Simply writing your changes to your iPod and exiting from `gtkpod` doesn't disconnect your iPod from your Ubuntu Linux system. Like any storage device on a Linux system, an iPod is mounted as part of your Linux filesystem so that you can access it as part of the standard directory hierarchy. By default, iPods are mounted on the directory `/media/ipod` when you attach them to your system. Although it is mounted as a part of your system's filesystem, an iPod's screen will display an international No symbol and a "Do not disconnect!" message. To safely disconnect an iPod from the filesystem after using `gtkpod`, you must first unmount it to guarantee that all disk updates have been saved and that the iPod's filesystem is marked as up to date and clean.

If an icon for your iPod is displayed on your desktop, right-click that icon and select Unmount from the context-sensitive menu that displays. If you do not see an icon on your desktop or selecting the menu command doesn't work for you, you can always resort to the command line. Linux provides the `eject` command to cleanly unmount a disk partition or mounted device. Before disconnecting your iPod from your system, you can execute the following command:

```
$ sudo eject /media/ipod
```

After you supply your password, this command will unmount the iPod filesystem, remove any desktop icons associated with the iPod, and will usually cause your iPod's screen to stop displaying the international No symbol and "Do not disconnect!" Message. You can then physically disconnect it from your system.

Working with CompactFlash and SD Cards

Though Ubuntu and Kubuntu can recognize a tremendous selection of digital audio players, digital cameras, and other devices, you may encounter a situation where your system doesn't recognize a specific device or that device doesn't provide a USB or FireWire port to enable you to connect it to your Ubuntu system. Luckily, most of these devices provide removable storage such as CompactFlash or Secure Digital (SD) memory cards that can be accessed from your Ubuntu or Kubuntu system using a generic USB media reader. These devices are quite common nowadays and can be purchased on eBay or at most camera stores. Using these devices with Ubuntu Linux is quite simple and makes it possible for you to copy your precious photos, files, or other information from just about anything to the safety of your Linux desktop computer system.

After ejecting the removable storage media from your camera or other device, insert it into the media reader, and attach that device to your system. If your system can identify the media as a storage device from a digital camera (in other words, if the flash card contains a directory called `DCIM`), it will give the same notificiation dialog that is used for digital cameras. In some cases, the notification dialog may simply offer to mount it because your system could not uniquely identify the storage device as having come from a camera. In that case, you can simply open the storage device in the suggested file manager and then copy the files off the media using that file manager. If the CompactFlash or SD card that you are exploring is from a digital camera, your photos will be located in subdirectories of the media's `DCIM` directory.

Once you are finished working with your removable media card, you should eject it by right-clicking on its desktop icon or file manager entry and selecting the Safely Unmount or Eject commands from the context-sensitive menu. These commands ensure that any pending changes are actually written to your storage card and that it is cleanly dismounted so that it can be safely disconnected from your system.

Summary

Portable consumer electronics devices like digital music players and digital cameras are essentially high-tech jewelry today. Almost everyone has at least one. Therefore, this chapter covered a lot of digital media ground, including how your system recognizes newly attached digital devices and how to customize the actions that it offers to perform, how to work with iPods and other digital audio players, and how to work with digital cameras or directly with removable CompactFlash or Secure Digital storage cards. Though I've been using Linux for over a decade now, the quality, power, and even (nowadays) the ease-of-use of its application ecosystem still impresses me. After reading this chapter, I hope that you feel the same way.

Chapter 22 discusses the sequence of events that occur when you boot or shut down your Ubuntu or Kubuntu system. The system startup/shutdown process is very different on Ubuntu and Kubuntu systems than on other Linux distributions because Ubuntu has innovated there to replace the traditional startup/shutdown system with a more modern, event-driven system. Understanding the system startup process is important to successfully integrate new services and capabilities into an existing system. Knowing where and how various services are started is also important when optimizing your system, because this enables you to easily disable specific services that you don't really need and that are, therefore, wasting CPU time on your system. Understanding the system shutdown process is important in terms of understanding how to cleanly shut down various services, in the right order, so that you avoid various system cleanup and housekeeping tasks the next time you boot your system.

Part III

Ubuntu and Kubuntu for System Administrators

IN THIS PART

Chapter 22:
Understanding the System Startup and Shutdown Processes

Chapter 23:
Adding, Removing, and Updating Software

Chapter 24:
Adding Hardware and Attaching Peripherals

Chapter 25:
Network Configuration and Security

Chapter 26:
Going Wireless

Chapter 22

Understanding the System Startup and Shutdown Processes

Nowadays, booting an Ubuntu or Kubuntu desktop system involves watching a friendly splash screen display a variety of system status messages and progress bars as it proceeds through the various commands that are part of the boot, initialization, and startup processes. This is all well and good, but sometimes it's nice to know what's actually going on under the hood. It's similarly interesting to know how to cleanly shut down your system when you want to turn it off for a while, and what it does during an orderly shutdown. (Describing what happens when you do a disorderly shutdown by pulling the plug is much easier—poof! The system's off, but the boot process next time will be a tad slower as the system checks for filesystem problems and repairs anything that it has to.)

This chapter describes the general phases of the Linux boot process and any things that are specific to Ubuntu and Kubuntu Linux. It then discusses how you can get detailed information about all of the processes executed when your system starts, so that you can look for ways to speed up your system's boot process, and discusses upcoming improvements that will further speed up the boot process in general. The chapter concludes with an explanation of what happens when you shut down your system, which is relatively short, because this is really just the reverse of the last stage of the boot process.

IN THIS CHAPTER

Understanding the Linux boot process

Monitoring the Linux boot process

Optimizing the Linux boot process

Understanding the Linux shutdown process

Overview of the Ubuntu/Kubuntu Startup Process

This section describes the process of booting Ubuntu or Kubuntu Linux, from the time that you turn your system on to the point at which your operating system of choice is running. Some of the low-level steps are x86-specific, but everything higher than the boot loader should apply to any hardware on which you can run Ubuntu or Kubuntu. (Welcome to the party, ARM fans! You can sit over with the SPARC and PPC holdouts for right now.)

The Boot Monitor or BIOS

When you first turn on a computer system, it loads a boot monitor or Basic Input/Output system (BIOS) from storage on the motherboard. This storage is usually a programmable, read-only memory chip (PROM) or a section of flash memory that is present on the board. The BIOS or boot monitor is a very low-level, hardware-oriented application that does some basic hardware initialization, performs some amount of hardware testing and verification (often optional), waits for keyboard or console input for some predetermined period of time, and then usually begins to execute a set of predetermined instructions to load another program into memory. These instructions load another program into memory from a predetermined location such as a portion of flash memory or the Master Boot Record (MBR) of a storage device such as a hard drive or CD-ROM, and then specify the memory address at which to begin executing the program once it has been loaded successfully.

On systems with limited hardware resources, such as many embedded systems, the program loaded into memory by a boot monitor is the Linux kernel, which is described in a subsequent section. However, the boot process is slightly more involved on most desktop and server computer systems. On these types of systems, an intermediate program is loaded into memory from the MBR that provides additional configuration and customization options for the boot process. On modern x86 systems, the program that is loaded into memory at this point is known as a *boot loader*, which is a configurable application that provides higher-level options for loading an operating system on the computer. The boot loader that is most commonly used on modern x86 systems (including 64-bit systems) is the Grand Unified Boot Loader (GRUB) although some Linux distributions still use an older, simpler boot loader known as the Linux Loader (LiLo).

The Boot Loader

Boot loaders typically consist of two stages: a minimal first-stage boot loader and a more elegant, second-stage boot loader. The size of a first-stage boot loader is constrained by the size of the MBR on a storage device, which is a single sector (512 bytes). The first-stage boot loader, therefore, primarily contains enough information to locate and load the second-stage boot loader, which is where all the action is.

As mentioned earlier, Linux offers two boot loaders, LiLo and GRUB. The LiLo boot loader was developed long before GRUB, and does not understand Linux filesystem formats—it identifies its second-stage boot loader (/boot/boot.b) and other boot-related information through direct offsets into the device where these files are stored. The GRUB boot loader, which is the boot loader used by Ubuntu, follows a similar process, but it uses an intermediate step in the boot process, first loading what GRUB calls a *stage 1.5* boot loader that understands a specific type of Linux filesystem and can, therefore, load GRUB's text-format configuration file directly from your machine's root filesystem. GRUB provides many different stage 1.5 boot loaders as follows:

- **e2fs_stage1_5**, for booting from a partition containing an ext2 or ext3 filesystem
- **fat_stage1_5**, for booting from a partition containing a DOS or Microsoft Windows VFAT filesystem
- **ffs_stage1_5**, for booting from a partition containing a Berkeley Fast filesystem
- **iso9660_stage1_5**, for booting from a partition containing an ISO9660 filesystem, such as a CD-ROM
- **jfs_stage1_5**, for booting from a partition containing a JFS filesystem
- **minix_stage1**, for booting from a partition containing a Minix filesystem
- **reiserfs_stage1_5**, for booting from a partition containing a ReiserFS filesystem

- **ufs2_stage1_5**, for booting from a partition containing a traditional UNIX filesystem
- **vstafs_stage1_5,** for booting from a partition containing a Plan 9 VSTa filesystem
- **xfs_stage1_5**, for booting from a partition containing an XFS filesystem

By default, Ubuntu and Kubuntu systems include stage 1.5 boot loaders for booting from ext2/ext3, JFS, Minix, ReiserFS, VFAT, and XFS filesystems.

The stage 1.5 boot loader then loads GRUB's second-stage boot loader and associated configuration files directly from the type of filesystem on which they are located.

NOTE **The fundamental difference between LiLo and GRUB is whether they load boot loader configuration and menu data from a filesystem, or whether the location of this information needs to be built into the boot loader. Though both boot loaders use a configuration file, you must rerun the LiLo command (`/sbin/lilo`) each time you modify LiLo's configuration file (`/etc/lilo.conf`), so that LiLo can repackage information about available menu options and kernels into the boot.b file that it locates by a direct filesystem offset, and so that it can update the location of this file in the MBR. GRUB's stage 1.5 boot loaders work around this problem by enabling GRUB to locate its configuration information and associated files in a filesystem. There is no need to rerun GRUB after modifying its configuration data.**

A second-stage boot loader typically provides access to a menu that enables you to choose between different ways of loading an operating system on your computer. Both GRUB and LiLo can be configured to display this menu for a short period of time, to display it only in response to specific keyboard input (such as pressing the Escape key), or to not display it at all. Once your system boots Linux, you can modify the configuration of the boot loader to automatically display a menu by modifying the boot loader's configuration file, which is either `/etc/lilo.conf` for LiLo or `/boot/grub/menu.lst` for GRUB.

NOTE **Ubuntu and Kubuntu Linux systems use the GRUB boot loader, so that's the primary focus of this section. For more information about LiLo (which is still used by some other Linux distributions, such as Xandros Linux), see `www.freshmeat.net/projects/lilo`, which is the FreshMeat project page for LILO.**

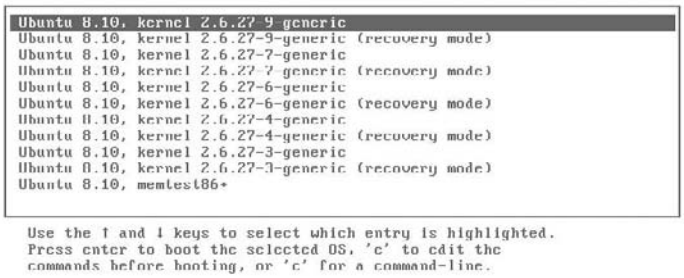
By default, the GRUB boot loader on Ubuntu and Kubuntu systems is configured to boot a default configuration without displaying a menu, but you can press the Escape key when GRUB first starts (do it quickly—you have five seconds) to display a menu that shows all of the available GRUB boot options. Figure 22-1 shows a sample menu from the GRUB boot loader on one of my Ubuntu systems. If you do not press Escape during the boot process, GRUB will automatically boot the default configuration that its configuration file directs it to use.

Some Linux distributions, such as Fedora Core, provide prettier GRUB menu screens than the one shown in Figure 22-1. My perspective on that is simple: "Who cares? I don't really want to deal with those screens anyway." You will have to interact with the GRUB menu if you set up a dual-boot machine and want to boot your other operating system for some odd reason, but other than that, GRUB boot menus should be something that you never need to see.

TIP **Each time that you install a newer version of the Linux kernel on your system (for example, through the Update Manager) the package installation process automatically updates your GRUB configuration to boot the new kernel and associated files. The GRUB menu entry for the previous kernel is still available on the GRUB boot menu. If you experience problems with a new kernel on a system, you can reboot a previous kernel by pressing Escape during the GRUB boot process, using the arrow keys to select the kernel that you want to boot and pressing Return.**

FIGURE 22-1

A sample GRUB boot menu



GRUB configuration files typically contain multiple definitions for ways to boot your system. Each set of related entries is referred to as a stanza, and looks like the following:

```
title          Ubuntu 8.10, kernel 2.6.27-9-generic
uuid           12de2e49-4579-452e-b079-14ba950773c0
kernel         /boot/vmlinuz-2.6.27-9-generic \
                 root=UUID=12de2e49-4579-452e-b079-14ba950773c0 \
                 ro quiet splash
initrd         /boot/initrd.img-2.6.27-9-generic
savedefault
boot
```

(The backslashes at the end of the kernel line and the next line are present for formatting purposes. This line and the next two would all be on a single line, with no backslash, in a normal menu.lst file.)

Each line in a stanza begins with a keyword that identifies the aspect of the boot process that it described. In the preceding example, these entries have the following meanings:

- title: Provides a label for this entry that is displayed on the GRUB menu.
- uuid: Identifies the Unique User ID that will be used to identify the absolute location of the runtime root filesystem partition in the subsequent `kernel` entry. Some systems (such as older Ubuntu and Kubuntu systems) simply identify the root filesystem, using GRUB's `root` keyword and an entry such as `root (hd0,0)` instead of using a UUID entry. In this example, the entry `(hd0,0)` means the first partition on the first hard drive.
- kernel: Identifies the kernel to boot, indicates how to identify the root filesystem, and provides kernel command-line boot parameters for that kernel.
- initrd: Identifies an initial RAM disk or initial RAM filesystem image that should be loaded by the kernel during the boot process.
- savedefault: Identifies this entry as the default GRUB boot entry.
- boot: Tells GRUB to boot using the configuration described in the current stanza.

These keywords are only a subset of the keywords supported by GRUB. For complete information about GRUB configuration files, see the online documentation for GRUB, which is provided in the grub-doc package, which you must install manually, using the Synaptic Package Manager or the apt-get application. Once you have installed this package, the complete GRUB documentation is available by typing **info grub** from any Ubuntu or Kubuntu command line. For even more information about GRUB, see the GRUB home page at `www.gnu.org/software/grub`. This page also discusses the new, improved, family-size version of GRUB, GRUB2, which is currently under development.

Loading the Kernel

A kernel specified in a GRUB stanza is typically a compressed kernel image that is loaded directly into memory. This compressed kernel is prefixed by some instructions that perform initial hardware and execution environment setup (stack definition, page table initialization, starting the swapper, and so on), uncompress the kernel image into high memory, store any initial RAM disk or initial RAM filesystem in memory for subsequent access, and then begin execution of the uncompressed kernel.

The kernel sets up interrupts, performs additional hardware initialization, and then uncompresses and mounts any initial RAM disk or filesystem that is found in the kernel image or is specified using GRUB's initrd keyword. If an initial RAM disk or filesystem is found or specified, the system follows the sequence of events described in the next section, "Loading and Using an Initial RAM Disk or RAM Filesystem." If no initial RAM disk or initial RAM filesystem is found in the kernel or identified using the initrd parameter, the kernel mounts the root filesystem identified on the command line, using the disk identifiers specified in GRUB's root entry, and then starts the standard system initialization process, as described later in this chapter in the section entitled "The Init Process."

Loading and Using an Initial RAM Disk or RAM Filesystem

Many Linux systems use a special type of filesystem as part of the boot process, before actually mounting the nontransient storage devices that are attached to your system. These initial filesystems are optional, and come in two different flavors: an *initrd* (initial RAM disk) or *initramfs* (initial RAM filesystem). On most desktop and server computer systems, these initial filesystems are typically used to load kernel modules, such as specialized device drivers that the kernel needs but that are not built into the kernel, and to perform specialized hardware and software initialization functions. On desktop and server systems such as Ubuntu and Kubuntu Linux systems, an initial RAM disk or initial RAM filesystem is almost always a transient part of the boot process. On embedded Linux systems with sufficient memory, the initramfs or initrd might be the actual root filesystem used when Linux is running.

Initial RAM disks are the traditional mechanism used by many 2.6 and all earlier Linux kernels to enable the kernel to load drivers and execute processes that must be available to boot successfully. Initial RAM disks are compressed files that provide an initial filesystem for use by the Linux kernel at boot time. Initial RAM disks are compressed files that contain an actual filesystem in a format such as ext2, romfs, and so on. Initial RAM disks are uncompressed into a pre-allocated portion of memory and are then mounted for use by the kernel. The primary problem with initial RAM disks is that they consume a substantial amount of system memory and, because they are actual block devices, require the kernel to contain drivers for whatever type of filesystem they use.

The Linux 2.6 kernel introduced the idea of initial RAM filesystems that exist directly in the kernel's file and directory cache. Initial RAM filesystems are identified or built into the kernel in the same way that initial RAM disks are, but consist of compressed archive files in the format produced by the UNIX/Linux cpio command. These archive files contain all of the files and directory entries for the initial RAM filesystem, and are directly unpacked into the kernel's file and directory entry cache, which is then mounted

as a filesystem. Using the file and directory cache as a filesystem substantially reduces the amount of memory required for the initial RAM filesystem because files and directories live in the kernel's page and directory entry (dentry) cache, where they would be copied anyway if they were used from a filesystem located on a block device. An initial RAM filesystem is an instance of the Linux tmpfs filesystem.

Linux systems that boot with an initial RAM disk execute the file /linuxrc from that filesystem after uncompressing and mounting the filesystem. This file is typically a command file, but can also be a version of the /sbin/init program that follows the standard initialization process described later in this chapter in the section entitled "The Init Process." (This is typically done when the initial RAM disk will be your runtime filesystem, as in embedded systems.) Linux systems that boot with an initial RAM filesystem execute the file /init from that filesystem after loading it into the file and directory entry cache. Like the /linuxrc file used with initial RAM disks, /init is typically a command file, but it can also be a version of the init program that follows the standard initialization process described in the next section.

Comparing Initial RAM Disks and RAM-Based Filesystems

Some of the more significant ways in which initial RAM disks and initial RAM filesystems differ are the following:

- You do not need any special privileges to build an initial RAM filesystem file or a kernel that contains one. Building an initial RAM disk typically requires root privileges unless you do so using an external application such as the genext2fs application.
- Initial RAM disks (initrd) are compressed files containing a filesystem in a specific format such as ext2, romfs, cramfs, squashfs, and so on. Initial RAM filesystems (initramfs) are compressed archive files in the format used by the Linux/UNIX cpio application.
- Initial RAM disks must be prepared outside the kernel and identified during the kernel build process if you want to link them directly into the kernel. Initial RAM filesystems can be prepared outside the kernel and identified during the kernel build process, but they can also be automatically created during the kernel build process. Versions 2.6 and later kernels enable you to identify a directory that contains the files that you want to have in your initial RAM filesystem or a configuration file that describes their contents.
- Initial RAM disks allocate and require a fixed amount of memory that is specified when building the kernel, regardless of the actual size of the files and directories that they contain. Any space not used by files and directories in the initial RAM disk is completely wasted. Initial RAM filesystems are directly allocated in the kernel's page and dentry cache and, therefore, only require allocating the amount of memory that is required by the files and directories that they contain.
- You can easily add a new script or file to an initial RAM filesystem that is specified as an external file in GRUB by using the gunzip application to uncompress the cpio archive, using the `cpio --append` command to add the new file to the archive, and using the `gzip` command to recompress the archive. You should only do this for testing purposes, because your modification will go away the next time you build a kernel or receive an update. You can permanently add the file or script to future initial RAM filesystems that you build by adding that file or script to the directory or script pointed to by the `CONFIG_INITRAMFS_SOURCE` kernel configuration variable for your kernel.
- Initial RAM disks automatically execute the file `/linuxrc` from the RAM disk. Initial RAM filesystems automatically execute the file `/init` from the RAM filesystem. The /init file in an initramfs filesystem runs with process ID 1, just like the init process in a standard Linux system.

The contents of an initial RAM disk or initial RAM filesystem differ based on the hardware, system capabilities, and type of Linux system that you are using. Desktop Linux systems typically use these filesystems to provide loadable device drivers or mandatory system setup procedures that are required to proceed with the boot process. Embedded Linux systems often use initial RAM disks or initial RAM filesystems as their actual runtime filesystems in the absence of usable disk, flash, or compact Flash storage. Both initial RAM disks and initial RAM filesystems provide only transient storage that does not enable you to save files across system restarts, but initial RAM filesystems are much more efficient and place fewer requirements on the kernel.

On desktop and server systems that use physical root filesystems, the last step in the /linuxrc or /init instructions is to mount the real root filesystem (identified in the kernel or using the root= kernel boot parameter), begin using that root filesystem as the new root filesystem (by using a pivot_root or chroot command), and start the init process, as described in the next section.

The Init Process

After loading the kernel and mounting a runtime root filesystem, traditional UNIX and Linux systems execute a system application known as the *init* (initialization) process, which is typically found in /sbin/init on Linux systems. Typically, the init process is process number 1 on the system, as shown in a process status listing produced using the ps command, and it is therefore the ancestor of all other processes on your system.

The traditional init binary used on most UNIX and Linux systems is part of a package of applications known as the SysVInit package. The SysVInit package uses a package of related command scripts known as the initscripts package, to identify the processes that are started when a system boots into different run levels and the order in which they are started. The SysVInit init program reads the file `/etc/inittab` to identify the way in which the system should boot and defines the other processes and programs that it should start.

As the name suggests, the SysVInit package and related processes have their conceptual roots in the Sys V version of UNIX, more properly UNIX System V, Release 4 (SVR4), which was released in 1989. Though elegant in its simplicity, the SysVInit package suffers from a lack of parallelism and flexibility, which can lead to longer system startup times than are absolutely necessary and can complicate the process of recognizing and using networked and removable devices. Various replacement init packages are available to replace SysVInit with a faster, increasingly parallel, and more responsive system initialization mechanism. The best known of these are initng (Init, the Next Generation, `www.initng.org`), Apple's launchd (`http://developer.apple.com/macosx/launchd.html`), Sun's Predictive Self-Healing architecture and Service Management Facilities (SMF) (`www.sun.com/bigadmin/content/selfheal`), and Upstart (`http://upstart.ubuntu.com`). Upstart was developed for the Ubuntu Linux distribution and is only used by that distribution (and Ubuntu-based distributions such as Kubuntu, Xubuntu, and so on, of course) at the time that this book was written, although it is scheduled to be adopted by other distributions (such as Fedora) soon. Upstart is designed to be portable and can easily be adopted by other distributions. Update is an event-driven system initialization package that increases parallelism in system startup tasks and provides a highly responsive framework for starting, stopping, and automatically restarting processes that depend on the availability of services that are started earlier in the system initialization process.

The next two sections explain how traditional Linux and UNIX systems can start different processes depending on how they were booted, and how Ubuntu's Upstart mechanism provides compatibility with existing mechanisms, while providing a more flexible, event-driven system startup procedure.

System Run Levels and Associated Command Scripts

Most desktop and server Linux systems use the concept of *run levels* to identify the applications that the system executes and the order in which they are executed. Different run levels are simply a way of defining a set of applications that the system administrator believes should be run when the system is being used in a certain way. For example, run level 1 is usually known as single-user mode, and is used for system maintenance. Only the core set of applications required for a system administrator to contact the system and perform maintenance tasks are running when the system is in run level 1. The run levels used by Ubuntu are the following:

- **0:** Halt
- **1:** Single-user
- **2:** Default multiuser run level on Ubuntu/Debian systems
- **3-5:** Other multiuser run levels, used for customized run levels
- **6:** Reboot

NOTE **Most Linux distributions other than Ubuntu and other Debian-inspired distributions use run level 3 as the default, nongraphical multiuser run level, and run level 5 as the default graphical multiuser run level.**

Run levels are a core concept of the SysVInit package and provide the traditional mechanism for system initialization on Linux systems. The concept of run levels is, therefore, well known and well understood by most system administrators. Even alternate system initialization and init packages, such as those introduced at the end of the previous section, are typically installed so that they emulate the spirit and behavior of the SysVInit package, while forsaking its internals. For example, Ubuntu's Upstart package installs itself in SysVInit compatibility mode so that it still follows the run-level convention.

The startup commands associated with each Linux run level are located in the directory `/etc/rcN.d`, where N is the number associated with a run level. The files in these directories have names of the form `SNNname` or `KNNname`, and are usually executable shell scripts (or, more commonly, are links to a common set of scripts) that contain lists of other system commands to execute. Files beginning with the letter "S" are executed when the system starts a specified run level. Entries beginning with the letter "K" are executed when the system leaves a specified run level. Commands are executed in the numerical order specified by the number NN. The "name" portion of an entry is user-defined but is generally evocative of the program or subsystem that the file starts.

Because Linux systems start many of the same commands at different run levels, the entries in the run-level command directories are usually just symbolic links to the master collection of system startup and shutdown scripts located in the directory `/etc/init.d`. The names of the symbolic links in the different run-level directories use the `SNNname` and `KNNname` naming conventions to impose a sequence on their execution. The names of the scripts in `/etc/init.d` that the run-level directories link to have the generic names. Using symbolic links rather than explicit startup and shutdown scripts for each run level makes it easy to add or remove services from a given run level without duplicating the master scripts, and leaves `/etc/init.d` as a single, central location where you can modify any of these scripts to propagate your changes to the startup/shutdown process for all of the associated run levels.

TIP **See "Optimizing the Boot Process," later in this chapter, for a discussion of some applications that can help you manage the startup/shutdown scripts and links associated with different Ubuntu run levels.**

Traditional SysVInit System Initialization

Systems that use the SysVInit and initscripts packages boot in the following sequence: Once the kernel is loaded and begins executing from its runtime root filesystem, it invokes a process known as the /sbin/init (initialization) process, which reads the file /etc/inittab to identify the basic processes and programs that it should start.

In the /etc/inittab file, the default run level for your system is identified with a line containing the string initdefault, as in the following example:

```
id:2:initdefault:
```

In this case, the default run level is 2. An additional entry in the /etc/inittab file identifies a command script that is run to initialize the system itself, as in the following example:

```
si::sysinit:/etc/init.d/rcS
```

NOTE **On Ubuntu and Kubuntu systems, the `/etc/init.d/rcS` command actually runs the `/etc/init.d/rc` command with an argument of `S`, which executes all of the scripts in the directory `/etc/rcS.d` that begin with an uppercase `S`, as described in the rest of this section. On some other Linux distributions, the system initialization script is a single script to which you manually add any initialization commands that you want to run.**

After the system initialization script is run, Linux distributions that use the SysVInit and initscripts packages then define what occurs when the system enters a specific run level. For example, a few lines later in the /etc/inittab file, you might see additional information about run level 2 in an entry like the following:

```
l2:2:wait:/etc/init.d/rc 2
```

This line specifies that when the system enters run level 2, the /etc/init.d/rc command will be executed with an argument of 2. This causes the rc (run command) command to execute all of the appropriate files located in the directory /etc/rc2.d. The files in this directory have names of the form SNNname or KNNname, and they are usually executable shell scripts (or are links to a common set of scripts) that contain lists of other system commands to execute. Files beginning with the letter S are executed when the system starts a specified run level. Entries beginning with the letter K are executed when the system leaves a specified run level. Commands are executed in the numerical order specified by the number NN. The "name" portion of an entry is user-defined but is generally evocative of the program or subsystem that the file starts.

Upstart System Initialization

- As introduced earlier, Upstart is an extremely interesting event-driven system initialization system that replaces the traditional init program and provides compatibility with the SysVInit/initscripts run-level mechanism through the `upstart-compat-sysv` package. Events are essentially string messages that can be sent by a process in response to a change in the state of something that it is monitoring. Event messages are sent once by a process. The control files for Upstart's version of `/sbin/init` is known as "job files" and are located in the directory `/etc/event.d`. Each of these are a text file that must contain at least the following:
 - One or more events in response to which that job file should perform some action. For example, the "start on startup" entry states that a job file should be executed when a startup event is received, and the "stop on run level 2" entry states that a job file should stop whenever a run level 2" event is received.

- An exec or script stanza that identifies the tasks to be performed in response to the events for which this job file is valid. An exec stanza is used to execute a specific binary and associated command line on the system. A script stanza provides shell script code that is executed by `/bin/sh` and must end with an "end script" statement.

Other keywords can also be used in Upstart job files in order to identify output devices, scripts to run before the primary exec or script section is executed, scripts to run after the primary exec or script sections completed, and so on. The pre- and post-scripts are intended to be able to initialize the environment required for script or exec commands, and to clean up or perform post-processing after they complete. Other Upstart commands are also available in job files. See `http://upstart.ubuntu.com/getting-started.html` for a complete list.

For example, the following is the Upstart job file `/etc/event.d/rcS`, which is used to emulate the behavior of a Linux system at startup time or in single-user mode:

```
start on startup
stop on run level
# Note: there can be no previous run level here, if we have one
# it's bad information (we enter rc1 not rcS for maintenance).
console output
script
        run level --set S >/dev/null || true
        PREVLEVEL=N
        RUNLEVEL=S
        export PREVLEVEL RUNLEVEL
        exec /etc/init.d/rcS
end script
```

- Two conceptual classes of jobs are managed by Upstart:
 - Tasks, which must complete and, therefore, essentially transition from stopped to started and back to stopped.
 - Services, which must be running and, therefore, essentially transition from stopped to started. Services have either "respawn" or "service" in their job file.

Upstart's init program generates the following system events:

- **startup:** Event issued when the system boots and `/sbin/init` first starts.
- **shutdown:** Event issued when the system is in the process of being shut down.
- **stalled:** Event issued when there are no jobs running and no pending events are in the queue.

Similarly, the shutdown tool provided with Upstart sends one of the following events when the system is being shut down, depending upon the way in which the shutdown command was issued:

- reboot
- halt
- poweroff
- maintenance (i.e., single-user mode)

You can also manually send a specific event from the shutdown command by invoking it as "shutdown -e event…."

- Upstart jobs also generate events themselves when they change state, making it easy to chain multiple job files together and guarantee that dependencies (such as the availability of required services or hardware/software initialization) are satisfied before changing state. Job files issue the following events:
 - **jobname/start:** Sent when a job is first started.
 - **jobname/started:** Sent once a job has actually started.
 - **jobname/stop:** Sent when a job is stopping.
 - **jobname/stopped:** Sent once a job has stopped.
 - **jobname:** For services, this event is generated once the associated task is running. For tasks, this event is generated once it has finished.

Processes on the system can send events through the Upstart control socket or by invoking the "initctl trigger EVENT" from within a script or job file. The `/sbin/initctl` utility enables a privileged user to communicate with the Upstart init daemon to start, stop, or monitor the status of Upstart jobs. When called with the list argument, the `initctl` command displays the status of all Upstart jobs, as in the following example:

```
$ sudo initctl listcontrol-alt-delete (stop) waiting
logd (stop) waiting
rc-default (stop) waiting
rc0 (stop) waiting
rc1 (stop) waiting
rc2 (stop) waiting
rc3 (stop) waiting
rc4 (stop) waiting
rc5 (stop) waiting
rc6 (stop) waiting
rcS (stop) waiting
rcS-sulogin (stop) waiting
sulogin (stop) waiting
tty1 (start) running, process 8513
tty2 (start) running, process 6485
tty3 (start) running, process 6488
tty4 (start) running, process 6478
tty5 (start) running, process 6479
tty6 (start) running, process 6493
```

The Upstart package also installs the `/sbin/start`, `/sbin/stop`, and `/sbin/status` commands, which are symbolic links to the `/sbin/initctl` binary and automatically invoke the `initctl` application to start, stop, or show the status of a specified Upstart job.

As mentioned earlier in this section, the Upstart system is installed on Ubuntu and Kubuntu systems in SysVInit emulation mode by default. When you start up a Ubuntu or Kubuntu system, Upstart's init process issues a startup event. This causes the script `/etc/event.d/rcS` to run, which executes the run-level command, triggering a run-level event. The script then execs itself again and stops when the run-level event is detected. Stopping this script triggers the `/etc/event.d/rc-default` script.

The `/etc/event.d/rc-default` script first checks if you booted the system with the "single" keyword on the kernel command line, which would indicate that you want to boot the system in single-user mode. If so, the script uses the `telinit` command to put the system in single-user mode. If not, the script checks for the presence of an `/etc/inittab` file. This file is not provided by default on Ubuntu 7.04 and greater systems, but this check gives you a way to override the system's default run level by creating this file with an appropriate `initdefault` entry (as described in the previous section). If this file is not found, the script automatically executes the run-level command to put the system in run level 2. This invokes the script `/etc/event.d/rc2`, which executes the scripts in the directory `/etc/rc2.d` in the standard order, and away we go!

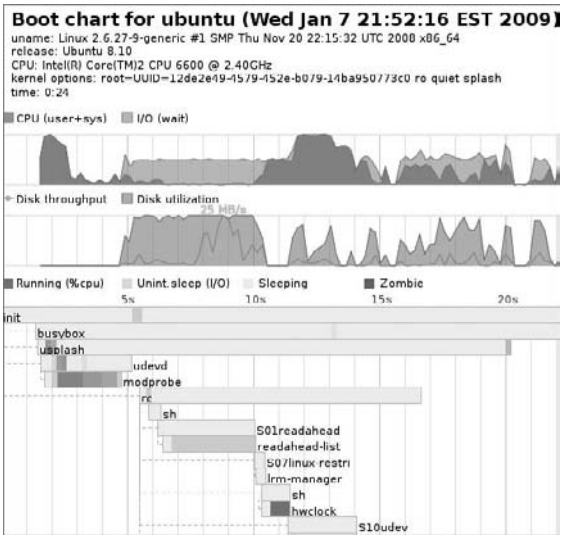
NOTE **Ubuntu's Upstart mechanism is a breath of fresh air that provides a faster, simpler, and more parallel boot mechanism while (at the moment) still preserving the core run-level concepts and associated startup script model that all Linux system administrators have become one with over the years. For more information about Upstart, see** `http://upstart.ubuntu.com`**.**

Examining the Boot Process with Bootchart

Bootchart (`www.bootchart.org`) by Ziga Mahkovec is an impressive and useful utility that monitors the boot process, creating data files that it uses to generate a graphical representation of your system's boot process, including detailed timing information. A portion of a Bootchart graphical output file is shown in Figure 22-2.

FIGURE 22-2

A section of Bootchart output



Bootchart isn't installed by default on Ubuntu or Kubuntu systems, but you can install it manually using either the Synaptic Package Manager or the apt-get utility. When installed on Ubuntu or Kubuntu systems, Bootchart adds itself to the initial RAM filesystem used by your kernel to capture data about the boot process and then adds a script to your runtime root filesystem's startup process (`/etc/init.d/stop-bootchart`, symlinked to `/etc/rc2.d/S99stop-bootchart`) that processes the collected data and generates the graphical record of the boot/startup process.

TIP **On Linux distributions that do not yet use an initial RAM filesystem (but still use an initial RAM disk), you can still use Bootchart by installing it on your system and modifying the GRUB boot entry for your kernel to include an `init=/sbin/bootchartd` entry so that the system runs the Bootchart data collection script before starting the `/sbin/init` program and executing the normal sequence of startup scripts.**

Bootchart generates Portable Network Graphics (PNG) graphics files in `/var/log/bootchart`. This file is named based on the day that it was created and a version number, so that multiple files created on the same day don't overwrite each other. You can configure Bootchart to produce graphics in SVG or EPS formats (or to preserve the data files that it creates) by modifying the `/etc/init.d/stop-bootchart` script.

Not only is it really cool to have a graphical record of the boot process, but that graphical record can be very valuable in terms of helping you identify startup scripts that are running that you don't want or need to run, extra invocations of system or X Window system processes, and so on. Bootchart adds very little overhead to your system's startup process, but the view of the system startup process that it provides can be invaluable.

TIP **Because Bootchart on Ubuntu and Kubuntu systems is added to the initial RAM filesystem, part of its installation procedure is to execute the update-initramfs command to update the initial RAM filesystem for your kernel. If a large number of kernels are listed in your /boot/grub/menu.lst file (or if you have manually created kernel entries in this file), make sure that the initial RAM filesytsem image that is being updated is the right one for your default kernel. If it is not, you can manually update the initial RAM filesystem for your kernel by using the sudo update-initramfs -u -k *KERNEL* command, where *KERNEL* is the kernel release that is returned by the uname -r command.**

Optimizing the Boot Process

Today's computers are faster than ever, with more memory and storage capacity than you'd find on the mainframes of yesteryear. Today's systems can easily run multiple applications at the same time and run more powerful software than ever before. At the same time, faster and better hardware has made most of us more impatient in terms of the amount of time that it takes for today's systems to boot.

The next sections discuss various ways of speeding up the boot process on your Ubuntu or Kubuntu system, ranging from the fairly obvious ("do less") to some whizzy technical mechanisms for optimizing how your system locates and accesses the applications and data that it needs during the boot process.

TIP **When performing system-level tasks like trying to optimize the boot process, it's hard to know what kind of improvement you've made unless you have an accurate record of before and after boot times. Before reducing startup scripts or adjusting readahead on your system, you may want to install Bootchart (described in the previous section) and boot your system with its original configuration in order to have Bootchart keep a record of what a system boot looked like and how long it took before you made your improvements. These Bootchart before and after records are also great conversation starters when trying to pick up a date at the local Internet café.**

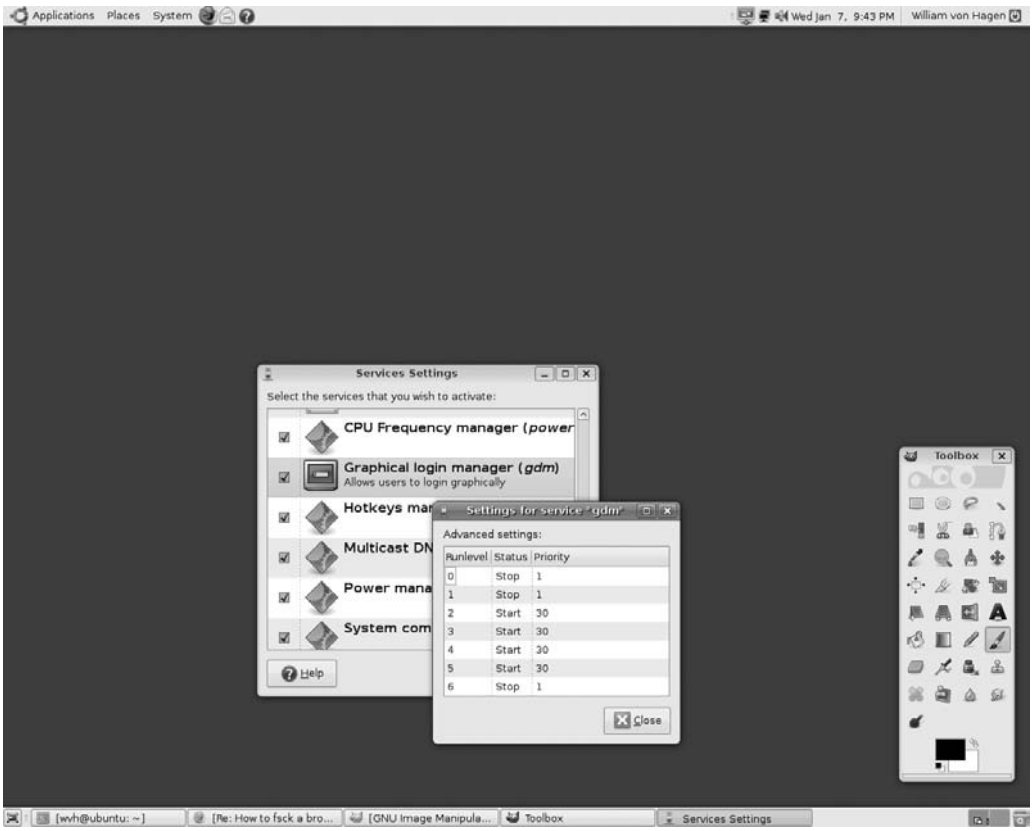
Minimizing the Amount of Software You Start

The section earlier in this chapter entitled "System Run Levels and Associated Command Scripts" identified the most common way of optimizing the boot process for a specific run level, which is by minimizing the number of startup scripts and associated commands that your Ubuntu or Kubuntu system executes when entering a specific run level. Though the Upstart system initialization mechanism increases the parallelism of your startup scripts, running fewer scripts at boot time reduces the number of processes that your system must execute before enabling users to log in to the system.

Ubuntu comes with an application (System ➪ Administration ➪ Services) that enables you to graphically control which services are started at various run levels when you boot your system. You can modify the run levels with which these services are associated or the sequence in which they are started by right-clicking on any entry, selecting the Properties context-sensitive menu, and left-clicking on the Status column entry for any run level that you want to modify. You can then select Start, Stop, or Ignore to specify the state of that service for that run level. You can only modify these values for services that are marked as active (i.e., ones that are preceded by a checkmark). Figure 22-3 shows this pop-up menu for the GDM (Graphical Display Manager) service at run level 2.

FIGURE 22-3

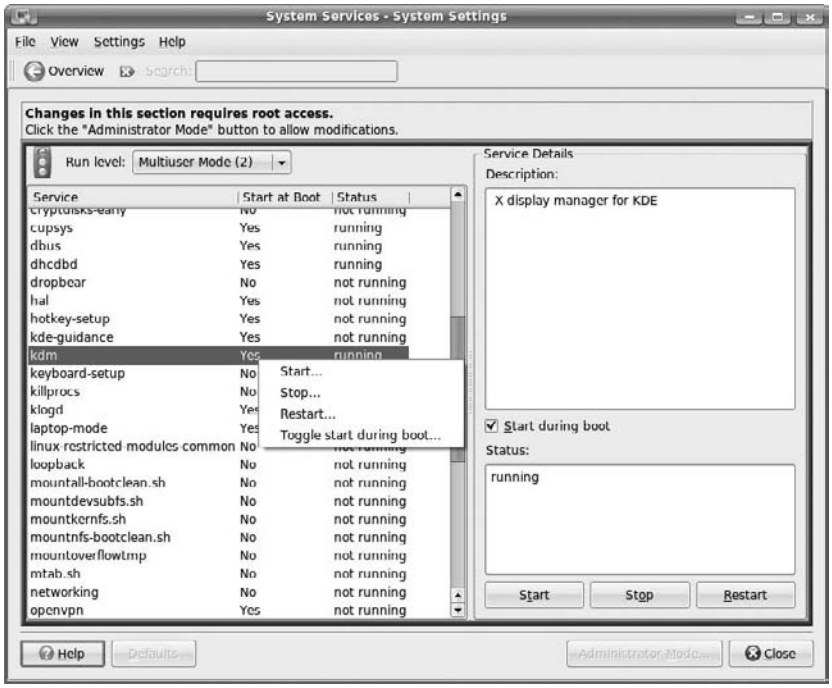
Configuring Services for Different Run Levels on Ubuntu



The KDE 3-based and KDE 4-based versions of Kubuntu provide a similar application, which is available by selecting the System Settings application's Advanced tab, and selecting System Services. The KDE 4 version of this application isn't complete, but the KDE 3 version of this application, shown in Figure 22-4, provides the same capabilities as the GNOME Services application.

FIGURE 22-4

Configuring Services for Different Run Levels on KDE3-based Kubuntu



Sadly, the disparity between many of the KDE 3 and KDE 4 applications is one of the reasons that I use KDE 4 relatively sparingly. Until the KDE 4 version of the System Services application catches up, you can always directly modify the system startup files that are associated with various services and their role in the boot process, as explained earlier in this chapter.

Optimizing Application Loading

While minimizing the processes that are running on your system is the primary (and, therefore, most popular) way of optimizing the Linux boot process, your system still has to load some number of executables in order to run the processes that it does need to run. If these executables are scattered all over the disk, your system has to position and reposition the heads of your disk drive(s) in order to locate, read, and load all of them. A significant amount of disk activity is normal at boot time, but reducing the

amount of disk activity (known as seeking) required to find all of the programs that you need to start can significantly reduce boot time.

Ubuntu and Kubuntu systems use a mechanism called *readahead* to attempt to minimize disk seeking. When you install a Ubuntu or Kubuntu system, it comes with a predefined set of applications and libraries that will be required when booting a standard Ubuntu or Kubuntu system. This list is organized based on the location of these files on your disk. When the system boots, it loads these programs and libraries into memory in the order that they're listed in the file, which should therefore minimize the disk activity required to load these files. The scripts and applications that need to access them during the boot process can subsequently access them more quickly, because these programs and libraries are preloaded into memory.

The readahead mechanism used by Ubuntu and Kubuntu systems stores the list of files and libraries required at boot time in the file `/etc/readahead/boot`, and stores the list of files and libraries required when starting a desktop in the file `/etc/readahead/desktop`. Since these files are pre-created at the Ubuntu/Kubuntu "factory," they can quickly diverge from the list of programs and associated libraries that your system actually uses at boot time, especially if you tend to add and remove applications that are started at boot time. Luckily, you can recreate these lists of files by doing what is known as a *profiled boot*, which simply involves booting your system with a special argument. When these files are recreated, they will accurately reflect exactly what your system needs during boot time, which can substantially reduce the time it takes to boot your system.

To do a profiled boot, do the following:

1. Boot or reboot your system and select your kernel from the GRUB menu. You may need to press the Escape key during the boot process in order to see the GRUB menu.
2. Once your kernel is selected, press the letter "e" to temporarily edit the boot stanza associated with that kernel. The boot stanza for the selected kernel appears.
3. Use the cursor keys to select the kernel line in that boot stanza and press the letter "e" again to edit that line. The kernel entry is displayed on a line by itself, with the cursor positioned at the end of the line.
4. Press the space key and type the word **profile**. Press Return to temporarily save the modified kernel boot line.
5. Press the "b" key to boot using the modified GRUB boot stanza. (The change that you have made is not saved to disk and will only take effect for the current boot sequence.)

Your system's boot process will take longer than usual, with more disk activity, since the system is tracking each application and library that is being loaded and its location on your disk. However, the next time that you boot your system, you should see some reduction in the time that it takes to boot your system. If you've loaded Bootchart, as suggested earlier in this section and discussed in "Examining the Boot Process with Bootchart," you will be able to compare the timing reports before and after, and you will hopefully see a significant reduction in boot time.

NOTE **You should repeat the profiling process after you make any major changes to your system's boot process, such as when installing applications with boot-time components, after significant system updates have occurred, and certainly after upgrading your system from one version of Ubuntu or Kubuntu to the next.**

The Ubuntu/Kubuntu Shutdown Process

To shut down your Ubuntu Linux system, you can select the System ➪ Administration ➪ Quit menu item, which displays the dialog shown in Figure 22-5.

FIGURE 22-5

Ubuntu's system shutdown dialog



Clicking Shut Down in this dialog begins the system shutdown process, which is essentially the opposite of the startup process. (You can click Cancel within 60 seconds to abort the shutdown process if you change your mind.)

To shut down a KDE 3-based Kubuntu system, select the K Menu ➪ Logout command, which displays the dialog shown in Figure 22-6.

FIGURE 22-6

The KDE 3-based Kubuntu 8.04 and earlier system shutdown dialog

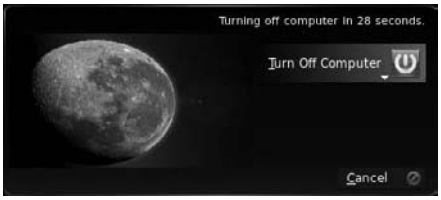


Selecting Turn Off from this dialog will immediately initiate the system shutdown process.

To shut down a KDE 4-based Kubuntu system, select the K Menu ➪ Leave menu, and select the Shutdown command to display the dialog shown in Figure 22-7.

FIGURE 22-7

The KDE 4-based Kubuntu 8.10 and later system shutdown dialog



Selecting Shutdown on this dialog begins the system shutdown process. (You can click Cancel within 30 seconds to abort the shutdown process if you change your mind.)

After selecting the appropriate item to shut down the type of system that you are using, the system logs you out, and the `/sbin/init` process begins to run the shutdown scripts associated with your current run level. When running SysVInit or (more probably) Upstart in SysVInit compatibility mode, these are the scripts beginning with the letter K in the `/etc/rcNN.d` directory, where NN is your current run level. The system then runs any shutdown scripts in the system initialization directory and turns off your system.

As a nod to x86-based PC tradition (and as a relatively quick silver bullet), the init process intercepts the traditional Ctrl+Alt+Delete key combination and can perform a specific action when this key combination is pressed. This key combination is known as the "three-finger salute" to many PC lovers.

On Ubuntu and Kubuntu systems, pressing the Ctrl+Alt+Delete key combination actually triggers an event, which is caught by the `/etc/event.d/control-alt-delete` job that was shown in the listing of Upstart jobs shown earlier in this chapter. When this event occurs, the system executes the following command:

```
/sbin/shutdown -r now
```

This command tells the system to run the `/sbin/shutdown` command with options that tell it to reboot after shutting down (-r), and to do everything now (now).

If you are running a command-line Ubuntu system, such as an Ubuntu server system, you can initiate the shutdown process from the command line by typing the following command:

```
/sbin/shutdown -r now
```

This command tells the system to run the `/sbin/shutdown` command with options that tell it to halt (i.e., turn off) after shutting down (-h) and to do everything now (now).

For more detailed information about the `/sbin/shutdown` command, see the online reference information available for this command by typing the man shutdown command from any xterm, GNOME Terminal, or Konsole window. You should probably do this at least once before actually shutting down your system—a few seconds isn't a lot of time to read a man page during the actual shutdown process.

Summary

This chapter dove down into the details of how Linux systems boot and the specific steps that your Ubuntu system follows when it boots. The Linux boot process is a very cool combination of hardwired and flexible steps that make it easy to customize the services that are started on your system and the order in which they are started.

Chapter 23 explores how to add, remove, and update software on your Ubuntu and Kubuntu systems. All Ubuntu and Kubuntu installations come with a great initial selection of software, but, as you've seen throughout this book, a huge amount of additional, useful software is readily available in the Ubuntu repositories. Chapter 23 explains how to retrieve that software, how to browse it, and how to figure out which software packages you need to install to find specific files or fulfill other requirements.

Chapter 23

Adding, Removing, and Updating Software

If you looked around the universe before settling on Ubuntu or Kubuntu Linux, one of the things that you were sure to notice is that "billions and billions" of Linux distributions are available, each with its own installer, favorite desktop environment and/or X Window system window manager, set of core applications, and some way of updating, expanding, and maintaining the software that comes with the distribution. When you come right down to it, the last item is the most significant aspect of a Linux distribution aside from the size and involvement of its user and development communities. There's actually an interesting loop between ease-of-use/ease-of-maintenance and the size of the user community. A Linux distribution that makes it easy to keep existing software up-to-date, install new software, and figure out what's on your system in the first place is a Linux distribution that more people are apt to use. (This assumes that someone is actually keeping the distribution up-to-date.)

As discussed in Chapter 1, I find the Ubuntu Linux community the most exciting, dynamic, and energetic Linux community that I've ever encountered. The Ubuntu community is a committed community of both users and developers. A vast pool of software is available for easy download and installation at any time. The Ubuntu and Kubuntu Add/Remove menu items make it easy to install new software from a preselected set of popular applications. If you want a bit more control or simply want to peek under the covers, the `dpkg`, `apt-get`, and `aptitude` tools make it easy to locate and install new software, and aptitude and the Synaptic and Adept Package Manager makes it easy to search for software that's relevant to your interests. Who could ask for anything more?

Updates and corrections to existing software are frequent, which is a good thing, and access to updated software is similarly easy. Both Ubuntu and Kubuntu run applets at startup (and regularly thereafter) that automatically let you know when updates are available to software that's installed on your system or the system itself, and keeping your system up-to-date requires three or four mouse clicks. Not too shabby!

IN THIS CHAPTER

- Adding/removing menu items
- Package management overview
- Using Ubuntu repositories
- Finding things on your system
- Using dpkg, apt-get, and aptitude
- Using the Synaptic Package Manager
- Using the Ubuntu Update Manager
- Using Kubuntu's Adept tools
- Cleaning up your system

The key to all of this simplicity and ease-of-use is the excellent DEB (for Debian) package format, developed for the Debian Linux distribution and therefore used by Ubuntu and Kubuntu. The DEB package format puts other software package formats, such as RPM, to shame—I personally would appreciate some electroshock therapy to forget the needless hours I've wasted trying to untangle twisty mazes of dependencies and requirements between various RPM packages. If you don't know what RPM is, you're lucky. Let's keep it that way.

NOTE **If you're familiar with RPM-based Linux distributions such as Fedora Core, Mandriva, Red Hat, or Yellow Dog, you're probably thinking, "Silly Human, yum solves most RPM problems!" Although this is probably true in a court of law, I don't feel that providing additional layers of syntactic sugar on top of a frustrating package management system such as RPM is the "right thing" or an actual solution. I prefer using smarter package formats such as DEB and smarter utilities such as the ones discussed in this chapter that were designed to do the right thing in the first place rather than continually shaving JoJo the dog-faced boy to make him look more human.**

TIP **If you must deal with packages that were created for other Linux distributions and are therefore in other package formats, see the section "Converting Packages from Other Package Formats" later in this chapter for information about converting package.**

The easiest way to install new software on an Ubuntu or Kubuntu system is by using the Add/Remove menu items found in the GNOME Applications menu or KDE K menu, respectively. These make it easy to install new software from a preselected set of popular applications, and are discussed in the first section of this chapter.

If you want more control over your system or will be doing system administration on multiple Ubuntu or Kubuntu systems, you will want to know about the underlying package management utilities used on those systems. The second section of this chapter provides an overview of the package management utilities used on Ubuntu and Kubuntu systems (and thus discussed in this chapter) so that you have a clear idea of "what does what" as you read through this chapter. The next section explains where the Ubuntu software that you can install on your system is located (in locations known as *repositories*) and how these repositories are organized, and provides some tips about other, unofficial repositories that you might find useful. Throughout this chapter, I'll be referring to the Ubuntu repositories, which are the repositories that house both Ubuntu and Kubuntu packages—there's no need to say "the Ubuntu and Kubuntu repositories" when referring to a single set of repositories. After discussing repositories, the hands-on sections of this chapter begin by explaining how to find out what's on your system and figure out what package provided specific files that are already on your system, and how to figure out what package you might want to install if you're looking for a specific file. This is most easily done from the command line, so that's our focus in this section.

Subsequent sections of this chapter focus on using the `apt-get` (command-line), `aptitude` (terminal-oriented), `synaptic` (fancy graphics for standard, GNOME-based Ubuntu), and `adept` (fancy graphics for Kubuntu) tools to locate and install new packages on your system. The `dpkg`, `apt-get`, `aptitude`, `synaptic`, and `adept` tools are knee-deep in online manual pages and Web sites that explain every conceivable combination of command-line options, so I'm not going to bore you by repeating those in this chapter. This chapter explains how to do common tasks using these tools, focusing on helping you get your work done.

The last two sections discuss how to use the Ubuntu Update Manager or the Adept Manager to keep your system up-to-date, and how to fine-tune your system by identifying and removing packages on your system that are no longer required because they were installed to satisfy dependency requirements for other packages that are no longer installed.

Installing Software with the Add/Remove Menu Item

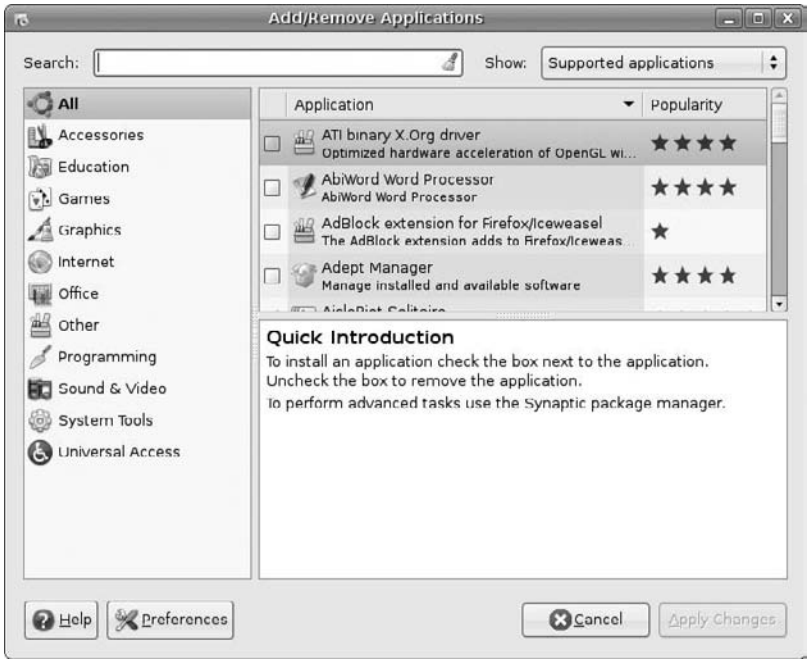
If you simply want to install a new software package on your Ubuntu or Kubuntu desktop system, the easiest way to do so is by using the Add/Remove menu item in Ubuntu's Applications menu, the Add/Remove Programs menu item in the K Menu on a KDE 3-based Kubuntu system, or the Add/Remove Software menu item on KDE 4-based Kubuntu systems. On Ubuntu systems, this menu item runs the `gnome-app-install` application; on Kubuntu systems it runs the Adept Installer. These menu items don't enable you to install every possible software package (for that, you will need to use tools such as `apt-get`, `aptitude`, `synaptic`, or `adept`), but do make it easy to install the most commonly used and popular packages on an Ubuntu or Kubuntu desktop system.

Using the Add/Remove Menu Item in Ubuntu

Selecting the Add/Remove item at the bottom of the Applications menu displays the dialog for the `gnome-app-install` application, shown in Figure 23-1.

FIGURE 23-1

Adding predefined applications to your Ubuntu desktop system



The elements of the Add/Remove Applications window are as follows:

- **Search field:** Enables you to search for a specific application by name if you aren't sure which category it falls under.
- **Show drop-down list:** Selecting this item displays a drop-down list that enables you specify the set of applications that should be displayed in the Applications pane. Your choices are as follows:
 - **All available applications:** Includes applications in the Applications window that are not officially supported by Canonical, Ltd, but are either available in the user-supported *Universe* and *Metaverse* sections of the Ubuntu repositories, or are third-party applications that are available through independent software vendors.
 - **All open source applications:** Includes all open source applications in the Applications window. This includes all applications that can be freely used, modified, and redistributed, including applications that are not officially supported by Canonical, Ltd, but are available in the user-supported Universe portion of the Ubuntu repositories.
 - **Canonical-maintained applications:** Shows only applications that are part of the official Ubuntu distribution and are thus officially supported by Canonical, Ltd. This is the default set of applications that is displayed in the Add/Remove Applications window.
 - **Third-party applications:** Shows only applications that are available through independent software vendors and are not part of Ubuntu.
 - **Installed applications:** Shows only applications that are currently installed on your Ubuntu system.

 For more information about Ubuntu repositories and how software is classified there, see "Ubuntu Repositories and Components" later in this chapter.
- **Software category pane:** Lists all of the predefined categories of software that the Add/Remove application knows about.
- **Applications pane:** Lists all of the applications that are available in the currently selected category, and their relative popularity (zero to five stars). The popularity index is designed to help you select which program to install from several similar choices.
- **Application description pane:** Displays a description of any currently selected software package.
- **Footer buttons:** You can perform the following actions:
 - **Help:** Displays detailed online help for the `gnome-app-install` application.
 - **Cancel:** Exits the application without making any pending changes to the Applications menu. Any changes that you've already applied using the Apply button are already done, and are not undone.
 - **Apply changes:** Updates the Applications menu based on your selections, installing any necessary software, and returning to the `gnome-app-install` application when done.

As an example, assume that you wanted to add the Gnumeric spreadsheet application to your Office menu. You would do the following:

1. Select the Add/Remove item from the Applications or K menu.

2. Select the Office category in the Software category List pane and browse the Applications pane for the `Gnumeric` application or enter **gnumeric** in the Search field and press Return.
3. Select the checkbox to the left of the Application pane entry for the Gnumeric spreadsheet.
4. Click Apply Changes to install Gnumeric.
5. Enter your password in the Administrative privileges dialog and press Return.
6. After Gnumeric is successfully installed, a dialog displays that provides buttons to enable you to Add/Remove additional applications or click OK to exit from the Add/Remove application.

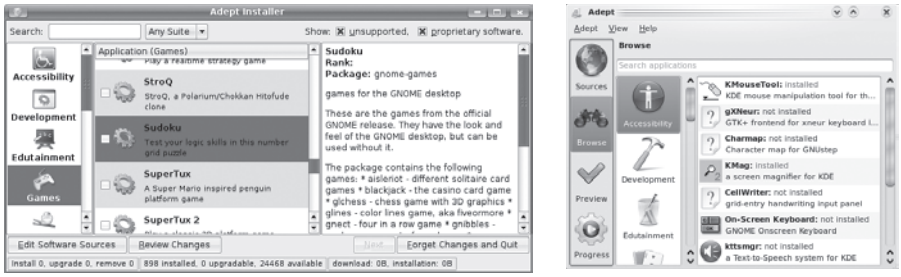
The Add/Remove applications automatically retrieve and install the requested software and any associated packages required to satisfy its requirements. If the software package you've requested conflicts with any existing software, a dialog displays that identifies the conflict and offers you the option to remove the conflicting package.

Using the Add/Remove Programs Menu Item in Kubuntu

If you are running Kubuntu, selecting the Add/Remove Programs item in the middle of a KDE 3-based Kubuntu system's K menu or the KDE 4-based Kubuntu system's Applications menu displays the dialog for the Adept Installer application, shown in Figure 23-2. The screen shown at left is the KDE 3-based Adept Installer screen, while the screen at right is the KDE 4-based Adept Installer.

FIGURE 23-2

Adding predefined applications to your Kubuntu desktop system



The portions of the Adept Installer window are as follows:

- **Header:** Provides general controls that help you locate packages. These are as follows:
 - **Search field:** Enables you to search for a specific application by name if you aren't sure which category it falls under. On KDE 3-based systems, you can use the Suite drop-down to restrict your searches to packages in the GNOME or KDE desktop suites. By default, searches look for matching packages in any desktop suite.
 - **Show checkboxes:** Only available in the KDE 3-based Adept Installer, these checkboxes enable you to prevent searches from looking for unsupported or proprietary packages by

de-selecting the associated checkbox. For more information about Ubuntu repositories and how software is classified there, see "Ubuntu Repositories and Components," later in this chapter.

- **Software category pane:** Lists all of the predefined types of software that the Add/Remove application knows about.
- **Applications pane:** Lists all of the applications that are available in the currently selected category.
- **Application description pane:** Available only on the KDE 3-based Adept Installer, this pane displays a description of any currently selected software package. In the KDE 4-based Adept Installer, selecting any pages in the Applications pane displays a description of that application. You must select an Install This Package checkbox below the description to subsequently install that package.
- **Buttons/icons:** Perform the following various actions in the KDE 3/KDE 4 Adept Installer:
 - **Edit software sources/Sources:** Displays a dialog that enables you to specify the repositories that will be searched for available packages. This application is discussed later in this chapter in the section entitled "Enabling Additional Repository Sources Using the Software Sources Tool."
 - **Review changes/Preview:** Displays a summary of any packages that have been selected for installation.
 - **Next/Apply changes:** (KDE 3 only) Displays Next if no packages have been selected for installation, and displays Apply Changes if some have.
 - **Forget changes and quit:** (KDE 3 only) Exits the application without installing any packages that are currently selected. Any changes that you've already applied using the Apply Changes button are already committed and are therefore not undone.
 - **Progress:** (KDE 4 only) Displays the status of any operations that are in progress.

As an example, assume that you wanted to add the KSpread spreadsheet application to your Office menu. You could do the following:

1. Select the Add/Remove Programs item from the Applications menu.
2. Enter your password in the Administrative privileges (kdesu) dialog and press Return.
3. Enter **spreadsheet** in the search box and press Return.
4. On a KDE 3-based system, select the checkbox to the left of the Application pane entry for the KSpread spreadsheet. On a KDE 4-based system, select the package listing to expand its description and select the Install This Package checkbox following the description.
5. Click Apply Changes (KDE 3) or Preview followed by Apply Changes (KDE 4) to install KSpread.
6. After KSpread is successfully installed, you can click the Return to Program Selection button to enable you to add or remove additional applications, or click Quit to exit from the Adept Installer application.

The Adept Installer application automatically retrieves and installs all requested software and any associated packages required to satisfy its requirements. If the software package you've requested conflicts with any existing software, a dialog displays that identifies the conflict and offers you the option to remove the conflicting package.

Package Management Software Overview

Ubuntu and Kubuntu Linux installations provide several different utilities for managing and querying software packages, ranging from low-level command-line tools to tools with sophisticated graphical interfaces. The following are the basic user-level tools provided as part of a standard Ubuntu or Kubuntu installation:

- `apt-get`: A command-line utility that provides subcommands, which enable you to install, remove, and manage packages on your system, both as individual packages and as components of a distribution.
- `aptitude`: A terminal-oriented utility that serves as a front end to lower-level utilities such as `apt-get` and `dpkg`. The `aptitude` utility provides both a quasi-graphical, menu and screen-oriented user interface and the ability to install, remove, query, and manage packages from the command line.
- `dpkg`: The basic command-line tool for installing, removing, querying, and generally managing packages. This utility uses even lower-level utilities, such as `dpkg-deb`, to perform package installation, removal, and manipulation, and `dpkg-query`, to search for packages, but these lower-level packages are not discussed in this chapter. Although handy to use directly .0001% of the time, the rest of the time, I find it more convenient to think of the capabilities provided by these low-level utilities as functionality that is provided by the `dpkg` utility itself.
- `dselect`: A terminal-oriented front end to the `dpkg` utility that provides quasi-graphical menus and an interactive display within the context of an xterm, GNOME Terminal, KDE Konsole, or other command-line environment. The `dselect` utility also accepts command-line options and arguments that you can use to initiate selected operations without selecting them from menus. This application isn't discussed in this chapter, because its functionality is available in other tools.
- `gnome-app-install`: Discussed earlier in this chapter, this is a graphical application that is started when you select Add/Remove Programs from the Applications menu on an Ubuntu system.
- `synaptic`: A graphical, X Window system tool for installing, removing, querying, and managing software packages on Ubuntu. This is my preferred package management utility for all day-to-day package installation operations.
- `update-manager`: A graphical, X Window system tool for identifying and installing updated versions of packages that are already installed on your Ubuntu Linux system.

Kubuntu systems do not provide the `synaptic` utility, but provide a set of graphical utilities that use KDE and Qt libraries rather than the GTK libraries used by synaptic. These applications, installed by default on Kubuntu systems, are as follows:

- `adept_installer` (KDE 3) or `adept installer` (KDE 4): A graphical, X Window system tool for identifying and installing new packages on your Kubuntu Linux system. On KDE 4 systems, executing the `adept installer` command only loads the portions of the Adept Package Manager that are relevant to installing new software packages. This is the Kubuntu equivalent of Ubuntu's `gnome-app-install` utility.
- `adept_manager` (KDE 3) or `adept` (KDE 4): A graphical, X Window system tool for installing, removing, querying, and managing software packages on Kubuntu. This is my preferred package management utility for all day-to-day package installation operations on Kubuntu systems. This is the Kubuntu equivalent of Ubuntu's `synaptic` utility.

- adept_updater (KDE 3) or `adept updater` (KDE 4): A graphical, X Window system tool for identifying and installing updated versions of packages that are already installed on your Kubuntu Linux system. On KDE 4 systems, executing the `adept updater` command loads only the portions of the Adept Package Manager that are relevant to updating software packages. This is the Kubuntu equivalent of Ubuntu's `update-manager` utility.

NOTE **The `dpkg`, `apt-get`, and `aptitude` utilities all use different databases to store information about installed packages and their state, and therefore all use different mechanisms for identifying and resolving conflicts and dependencies between software packages. (The synaptic and Adept utilities use the same database as the aptitude utility.) The `dpkg` utility stores its information about installed and available packages in files and directories under `/var/lib/dpkg`. The `apt-get` utility stores the information that it uses in files and directories under the directory `/var/lib/apt`. The `aptitude` utility stores the information that it uses in files and directories under the directory `/var/lib/aptitude`. When using either `apt-get` or `aptitude`, it is important that you run either the apt-get `update` or aptitude `update` commands in order to ensure that the package state and availability information used by these applications is up-to-date. Similarly, if you want to experiment with `dselect` after having installed packages using `apt-get`, `aptitude`, or `synaptic`, you must use the `dselect update` command before using `dselect`, to ensure that `dselect` (and therefore `dpkg`) is aware of the state of all installed and available software.**

The `apt-get`, `aptitude`, Adept, and Synaptic Package Manager utilities also depend on the information about available repositories and associated packages maintained in a storage area which is generically known as the "apt cache." The apt cache actually consists of several files that, by default, are located in the directory `/var/cache/apt` on your Ubuntu or Kubuntu Linux system.

This chapter focuses on discussing the `apt-get`, `aptitude`, `dpkg`, `synaptic`, `update-manager` tools in the context of performing common tasks that each is best suited for, based on my experiences with them. If you are using Kubuntu, the `apt-get`, `aptitude`, and `dpkg` tools are available for command-line use, but you should focus on the discussion of the Adept tools for information about graphical tools. As a fan of graphical tools that "do the right thing," most of the examples of installing packages throughout this book use the Synaptic Package Manager or Adept Manager tools, which I find both powerful and easy to use. In general, however, the mantra for this chapter is "use the right tool for the job."

All of the automated aspects of the package management utilities available on Ubuntu and Kubuntu systems rely on obtaining packages from the Ubuntu software repositories. The next section explains software repositories, how they are organized, the system configuration file that identifies them, and how to work with this file to ensure that your system has access to all of the latest and greatest software that is available for it.

Ubuntu Repositories and Components

A repository is exactly what the name suggests, a storage site for objects of some sort. In the case of the Ubuntu repositories, the objects stored therein are all of the source and binary packages that make up a variety of Ubuntu Linux distributions, including both standard GNOME-based Ubuntu Linux and the KDE-based Kubuntu Linux. When you download an ISO image (an image of a CD or DVD that conforms to the International Standards Organization 9660 standard for CD-ROM filesystem), the CD that you burn from this image and then use to install Ubuntu or Kubuntu Linux provides the CD boot environment (known as *Isolinux*), an installer and the applications required to support it, a basic set of packages, and some documentation. As part of the install process, your system retrieves the majority of the Ubuntu and Kubuntu Linux distributions from the Ubuntu repository over the Internet.

NOTE **Because the packages for both Ubuntu and Kubuntu are stored in the same repositories at Canonical, these repositories are generally referred to as the *Ubuntu Repositories* throughout the remainder of this section. This isn't intended to slight Kubuntu, but rather to demonstrate that the central Canonical repositories contain the entire spectrum of officially and user-supported software for both Ubuntu and Kubuntu.**

Making ISO images of Linux distributions available over the Internet is nothing new—it's the standard way of distributing most Linux distributions nowadays. Putting the Ubuntu repositories directly on the Net has two primary advantages for Ubuntu and Kubuntu fans:

- Keeps the installation media down to a single CD and ISO image
- Makes every part of Ubuntu and Kubuntu Linux instantly available to any computer that is connected to the Internet, anywhere, including the latest updates and additions

To organize its repositories along lines that are important to the Ubuntu and Kubuntu folks, to many individuals, and to many corporations, the Ubuntu repositories are organized into four components, which are basically ways of classifying Linux software along support and licensing guidelines. The Ubuntu repositories contain four basic components:

- **main:** The main component contains binary packages (and source packages, in most cases) for Linux software that is officially supported by Canonical, Ltd., and can be freely redistributed. This does not mean that everything in the main repository component is GPL—the main repository component can include software that is distributed in binary format, such as firmware and binary fonts. The main portion of the repository is designed to provide everything that most people will need for a fully functional Linux desktop or server system that is fully supportable by Canonical, Ltd.
- **restricted:** The restricted component contains source and binary packages for commonly used software that is not available under a completely free license. Packages in the restricted repository component are not guaranteed to be completely supportable by Canonical, Ltd., but are provided because they are necessary to use Linux on certain hardware. For this reason, some items from the "restricted" repository are included on Ubuntu and Kubuntu installer CDs in subdirectories of `dists/`*`dist-name`*`/restricted` and `pool/restricted`. Restricted items typically include drivers for specific network interface cards (NICs) and video cards.
- **universe:** The universe repository component contains binary and source packages for the rest of the free and open source software that is commonly associated with Linux systems, but which is available under a variety of different licenses and is not guaranteed to be supported by Canonical, Ltd. There's simply too much software in the universe (pardon the expression) for Canonical to guarantee support for all of the software that it contains
- **multiverse:** The multiverse repository component contains binary and source packages for software that is not released under a license that meets the Ubuntu guidelines for free and open source software (`www.ubuntu.com/ubuntu/licensing`). This software is not supported by Canonical, Ltd., and it is your responsibility to verify that you satisfy the requirements imposed by software from the multiverse component.

NOTE **The official descriptions of the official Ubuntu repositories are located at `www.ubuntu.com/ubuntu/components`—I've tried to distill them down to their essentials, but if you're a licensing fan, there's plenty of good reading there.**

Two additional repository components also exist, known as the *backports* and *partner* repositories:

- The **backports** repository component contains updated or bleeding-edge versions of Linux software, including software that may already have an "official" version in one of the other repository components. Software in this repository component is not supported by Canonical but is provided because newer versions of existing software packages often provide some critical feature that users may require.
- The **partner** repository component contains software that is produced by various partners of Canonical. Software in this repository component is not supported by Canonical but is provided as a convenience to Ubuntu and Kubuntu users who need access to software from Canonical's partner companies.



If you are curious about the packages that are installed or available in the Ubuntu repositories for different Ubuntu releases, see the Web page at `http://packages.ubuntu.com`. This page enables you to select the distribution that you are interested in, and then view a list of all of the packages, organized into various logical categories. It also provides a search capability so that you can search for specific packages.

Enabling Additional Repository Components

The most common example of a situation in which you will want to access repository components other than the main and restricted components is when working with audio and video applications. The wide variety of CODECs (compressor/decompressor modules) used to encode digital audio and video, the platform-specific roots of many of these CODECs, and the hoops that many media companies make you jump through in order to actually play many digital audio and video formats on an Ubuntu Linux system make it necessary to push the boundaries of free software licensing.

By default, the main, restricted, universe, and multiverse components of the Ubuntu repositories are enabled when you install an Ubuntu or Kubuntu Linux distribution. "Enabled" means that the online sources of these repository components are included in the default `/etc/apt/sources.list file`, which is the file that all of the package installation and management utilities discussed in this chapter consult when looking for new and updated software. This file is a text file that contains entries, which describe the location and name of different online repositories and the repository components that are available there. The general format of an entry in the `/etc/apt/sources.list` file is as follows:

```
deb  URI distribution component-or-package(s)
```

The `deb` field identifies the package format; the URI is a universal resource identifier for the location where the repository can be found, and is typically an FTP or HTTP URL. The distribution field is typically the name of a standard Ubuntu distribution, but can be anything that is used to identify a subdirectory at the URI where associated packages can be found, and the component-or-package is one (or more) names of repository components or packages that are also used to navigate the directory tree at the URI and locate available packages. If multiple component-or-package names are provided, each identifies a different directory at the URI.

The entries in the `/etc/apt/sources.list` file for the main and restricted repository components of the 8.10 release of Ubuntu Linux, for example, are as follows:

```
deb http://archive.ubuntu.com/ubuntu/ intrepid main restricted
deb-src http://archive.ubuntu.com/ubuntu/ intrepid main restricted
```

The first of these identifies the source of binary packages, whereas the second identifies the source of source packages. A similar pair of entries provides access to updated versions of the packages in these repository components:

```
deb http://archive.ubuntu.com/ubuntu/ intrpid-updates main \
   restricted
deb-src http://archive.ubuntu.com/ubuntu/ hardy-updates main \
   restricted
```

The latter two entries are broken across two lines for formatting purposes in this book; they appear on single lines in the actual `/etc/apt/sources.list` file.

NOTE **The examples in this section reflect the `/etc/apt/sources.list` settings for the Ubuntu and Kubuntu 8.10 (Intrepid) release. If you are working with a different release, you should replace the word "intrepid" in these examples with the name of the Ubuntu/Kubuntu release that you are using, such as gutsy (the Ubuntu 7.10 release), hardy (the Ubuntu 8.04 release), or the name of any release that follows the 8.10 release (such as jackalope, the 9.04 release). You should not mix repository entries for different releases in an `/etc/apt/sources.list` file because the repositories for a newer release will typically provide the same software as the repositories for older releases, and will be compiled against a set of system libraries that are found on the newer release. See the section "Mixing Ubuntu and Debian Repositories" later in this chapter, for information about some cases in which you may want to have entries in your `sources.list` file that are not consistent with the Ubuntu/Kubuntu release that you are using.**

Finally, a similar pair of entries provides access to security fixes for the packages in these repository components:

```
deb http://security.ubuntu.com/ubuntu intrepid-security main \
    restricted
deb-src http://security.ubuntu.com/ubuntu intrepid-security \
    main restricted
```

Lines beginning with a hash mark (#) in the `/etc/apt/sources.list` file are comments. If you look through the entries in the default `/etc/apt/sources.list` file on your system, you will note that entries for access to the partner repository are already present in this file, but are commented out. You can enable access to this repository by using `sudo` to edit this file in your favorite text editor. To do this, edit the `/etc/apt/sources.list` file using your favorite text editor, using a command like the following:

```
$ sudo emacs /etc/apt/sources.list
```

Remove the hash marks at the beginning of the lines that enable the partner repositories and save the updated file. The updated entries in a sample `sources.list` file would look like the following:

```
deb http://archive.ubuntu.com/ubuntu/ intrepid partner
deb-src http://archive.ubuntu.com/ubuntu/ intrepid partner
```

You could, of course, simply add these entries to the end of the lines that already provide access to the main, restricted, universe, or multiverse components of the Ubuntu repositories, but I prefer to keep them separate just for clarity regarding licensing differences—and because that stays closest to the original format of the default `sources.list` file.

NOTE **Just to be perfectly clear, installing software from the universe and multiverse repositories may add software to your system that has licensing requirements that you or the company that you work for find odious (such as MP3-related, or video-related software), for which support is not guaranteed by the Ubuntu folks (even if you purchase support, there are no commitments to support it), and which may be difficult to get support for anywhere. If you are configuring systems for enterprise use, you should make sure that your firm is comfortable with the licenses provided with non-free software in the universe and multiverse repositories. I use all of the software discussed in this book, and am therefore my very own test case for everything this book discusses, but your mileage (and your responsibilities) may vary.**

Other Repositories of Interest

As you can see from the examples in the previous section, the standard Ubuntu repositories are all hosted on systems that are in the ubuntu.com domain. This makes perfect sense because this is the only way that the Canonical folks can guarantee the integrity of their repository and the adherence of various components and their contents to the associated licensing guidelines. However, additional repositories of Ubuntu software are also available on the Internet (I know that's a surprise), maintained by various groups and individuals. These alternate repository sites may simply mirror the contents of the official Ubuntu repositories to "share the load," or may provide software that is not found in the official Ubuntu repositories. In the latter case, this software may provide even newer versions of popular software than is found in the backports repository component, or may provide up-to-date versions of locally maintained software. As mentioned earlier, if you are interested in working with audio and video on your Ubuntu system, you may find that these alternate repository sites provide versions of related software with bleeding-edge features that you need. In general, such sites will provide examples of the entries that you should add to your `/etc/apt/sources.list` file in order to be able to list and retrieve software from them. This is typically done by concatenating new entries to the end of this file using a text editor, but can also be done using the Software Sources tool, as discussed in the next section.

Unofficial, but useful, software packages are one thing, but a more common need is to locate and install software with various licensing restrictions. A classic Linux "problem" is the fact that many of the CODECs and much of the related software used to play CDs, audio files, DVDs, and video files is proprietary software that requires a software license, is unavailable for Linux systems, and sometimes both. This is not the fault of Linux, but is rather a side-effect of greed and cluelessness in the software and audio/video media industries, as well as location-specific laws regarding software patents, copyrights, and intellectual property. Even though the last few Ubuntu and Kubuntu releases have provided increasingly easy ways of retrieving problematic audio/video packages, you may still need to manually retrieve some of the packages that you need from unofficial (or certainly non-Canonical) repositories. One good site for information about such packages is the Medibuntu (Multimedia, Entertainment & Distractions In Ubuntu) site at `www.medibuntu.org`. The Repository How To link there points to documentation on how to add repositories to your `/etc/apt/sources.list` file that will enable you to play your legally purchased audio and video CDs and DVDs on your Ubuntu or Kubuntu system.

Enabling Additional Repository Sources Using the Software Sources Tool

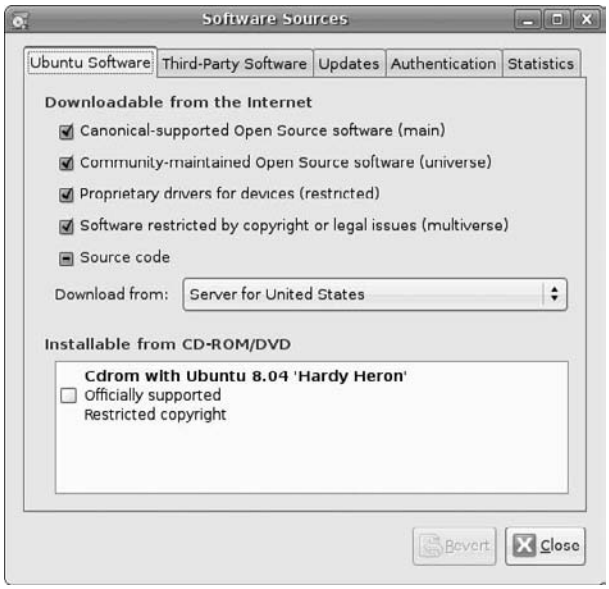
Editing a text-format configuration file is pretty old school Linux, and is the sort of thing that you may not want to mention to your Windows or Mac OS X friends. (You'll still want to be able to do it, because it's fast and easy, but we'll keep that sort of thing our little Linux secret, okay?) For fans of graphical

interfaces, both Ubuntu and Kubuntu provide a Software Sources tool that gives you a graphical view of the entries in your `/etc/apt/sources.list` file and makes it easy for your to enable, disable, or add entries with a few mouse clicks.

To start the Software Sources utility on a Ubuntu system, select the System ➪ Administration ➪ Software Sources menu item. After supplying your password to enable access to this administrative tool, the Software Sources tool's main dialog displays, as shown in Figure 23-3.

FIGURE 23-3

The main Software Sources dialog on Ubuntu



On Kubuntu systems, you can start the equivalent application from the Adept Installer by clicking the Edit Software Sources button (KDE 3) or the Sources icon (KDE 4), or from within Adept itself on KDE 3-based systems by selecting the Adept menu's Manage Repositories menu command. If you're really a command-line fan, you can also start this application from the command line on a Kubuntu system by executing the following command:

```
kdesu /usr/bin/software-properties-kde
```

The only real differences between the Ubuntu and Kubuntu versions of this application are that various tabs, menus, and available sources (such as the installation CD) may say Kubuntu rather than Ubuntu. Figure 23-4 shows this same dialog on a Kubuntu system. The rest of the screenshots in this section are from an Ubuntu system, but that shouldn't confuse anyone.

FIGURE 23-4

The main Software Sources dialog on Kubuntu



As you can see in Figures 23-3 and 23-4, the main software preferences dialog displays all of the entries in your current `/etc/apt/sources.list` file, regardless of whether they are commented out or not. Entries that are commented out in the `/etc/apt/sources.list` file are present, but the checkbox beside their name is not selected. This gives you a convenient overview of all of the valid, standard repository entries in your `/etc/apt/sources.list` file.

The Software Sources dialog also makes it easy for you to add third-party repositories (also often referred to as "custom" repositories to the repositories that the Ubuntu and Kubuntu software management tools will search for available and updated packages. This is done via the Software Sources application's Third-Party Software tab, shown in Figure 23-5.

To add a new repository or repository component to the list of repositories that your system will search for software packages, select the Third-Party Software tab and click Add. This displays the dialog shown in Figure 23-6.

Although not particularly elegant, this dialog makes it easy for you to cut and paste a repository identification entry from a mail message or browser session where you found it, and automatically add it to your `/etc/apt/sources.list` file. To return to the dialog shown in Figure 23-5 without adding a custom repository entry, click Cancel. To save a new entry after typing it in or cutting and pasting, click Add Source, which adds the new entry and returns you to the dialog shown in Figure 23-6.

In addition to the Ubuntu/Kubuntu Software and Third-Party Software tabs, the Software Sources tool provides three other tabs that let you manage different aspects of your repositories and the update process. These tabs enable you to manage different aspects of system updates, how your system identifies valid repositories, and also enable you to provide some general statistical information about recent updates to your system to help the Ubuntu folks improve the repository configuration and software update tools.

FIGURE 23-5

Adding third-party repositories in the Software Sources tool



FIGURE 23-6

Adding a specific third-party repository

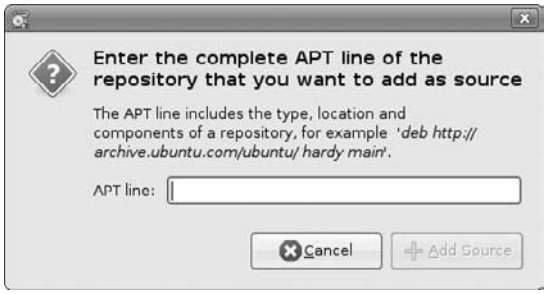


Figure 23-7 shows the Software Sources tool's Updates tab, which enables you to specify when your system checks for updates and how it responds to any available updates that it detects.

By default, your system automatically checks for updates on a daily basis, displaying an "Available Updates" icon in the GNOME panel at the top of your screen (or the KDE panel at the bottom of your screen if you are running Kubuntu) when updates are available. If you are using a laptop or other system that only connects to the Internet infrequently, you may want to decrease the frequency with which your

system checks for updates by selecting another value from the Check for updates automatically entry's drop-down menu, or disable this item entirely. If you disable this item, you can always check for updates manually using Ubuntu's Update Manager, which is discussed later in this chapter in the section "Using the Ubuntu Update Manager" or the Adept Updater if you are running Kubuntu. If you are performing regular checks for updates and are using a system that is often (or always) connected to the Internet, you may want to select the Download updates in the background, but do not install them to make sure that packages are already available on your local system when you actually execute the update process. Similarly, you may want to select the Install security updates without confirmation option if you want to automatically retrieve security updates and install them so that your system is always secure against any known attacks or software exploits.

FIGURE 23-7

The Updates tab in the Software Sources tool

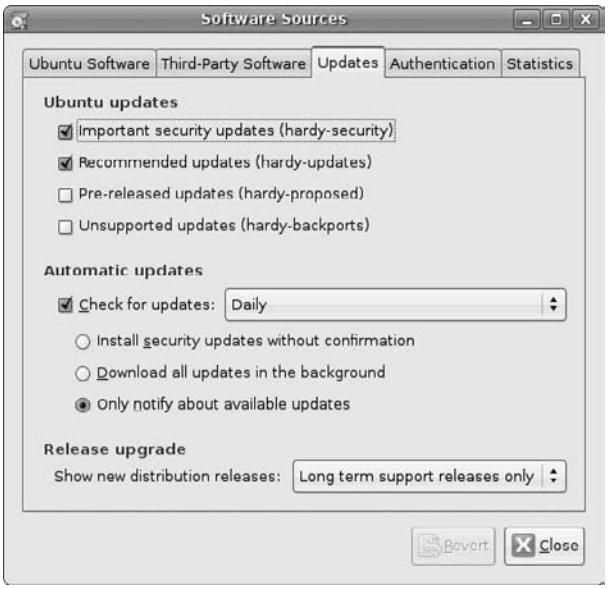


Figure 23-8 shows the Software Preferences tool's Authentication tab, which simplifies the process of importing identification keys for new repositories into the keyring used by Ubuntu's package management utilities.

Each package in any repository is signed with the key for that repository to verify its authenticity and validity. The Ubuntu package management tools will not retrieve packages from unauthorized or unidentifiable repositories. Adding the authentication key that identifies a nonstandard repository as a valid package source requires several steps from the command line, using hard-to-remember options to the `gpg` command in conjunction with the `apt-key` utility. Clicking the Import Key File button enables you to navigate to and select any file that contains an ASCII representation of a PGP authentication key and import it for use by Ubuntu's package management tools.

FIGURE 23-8

The Authentication tab in the Software Sources tool



After making any changes that you desire to the list of repositories available to your system, the way in which system updates are performed, or the Authentication keys available to the Ubuntu and Kubuntu package management utilities, click Close to exit the Software Sources application. If you have made any changes, the Software Sources tool will display a dialog informing you that your channel (repository) information is out-of-date, as shown in Figure 23-9.

FIGURE 23-9

Change notification when exiting the Software Sources tool



To automatically update the cache information used by the Ubuntu and Kubuntu package management utilities, click Reload. To exit without updating the cache, click Close. If you exit without reloading the cache information, you will need to run `apt-get update` or `aptitude update`, or use the Check button in Ubuntu's Update Manager, the Reload toolbar icon in the Synaptic Package Manager, or the Sources icon in the Adept tools (KDE 4) before you are guaranteed that your system is aware of all available packages in all valid repositories.

Problems Adding or Accessing Nonstandard Repositories

After modifying your system's `/etc/apt/sources.list` file, you may find that packages from new repositories that you've added aren't actually showing up in the list of available packages. First, make sure that you've updated the list of available repositories and packages, as described at the end of the previous section. This should identify any inaccessible repositories and outright syntax errors in the `sources.list` file.

Next, when running `apt-get update` to update your system's idea of repository contents, seeing messages like the following means that your system does not have the right authentication keys to add packages from a specified repository to its cache:

```
Ign http://ftp.us.debian.org sid Release
Ign http://ftp.us.debian.org unstable Release
```

Similarly, you may see messages at the end of the output from `apt-get update`. These are related to the previous messages, but are somewhat less subtle because they explicitly identify missing keys and associated repositories.

```
W: GPG error: http://ftp.us.debian.org testing Release: \
     The following signatures couldn't be verified because the \
     public key is not available: \
     NO_PUBKEY 010908312D230C5F
W: GPG error: http://ftp.us.debian.org unstable Release: \
     The following signatures couldn't be verified because the \
     public key is not available: \
     NO_PUBKEY 010908312D230C5F
```

In these cases, the associated entries in `/etc/apt/sources.list` point to repositories whose keys are not present in the keyring used by the Ubuntu package management utilities. See the man page for the `apt-key` application for information on retrieving and installing keys into this keyring.

Exploring Your System Using dpkg and Friends

Package management software not only simplifies software installation, removal, and updates, but also provides great opportunities for asking questions about the software that is installed, or could be installed, on your system. Some of the command-line package management utilities that are installed on your system make it easy to find out exactly what software is installed on your system, what a specific package contains, what packages provided certain files, and so on. The next few sections describe how to answer some basic questions about existing files and available packages on your Ubuntu system.

Listing the Packages That Are Installed on Your System

The `dpkg` command's `-l` option provides the easiest way of listing the packages that are installed on your system. (This option can also be specified in verbose option format as `--long`.) You can execute this command with no other argument to produce a (long) listing of every package that is installed on your system, which looks something like the following:

```
$ dpkg -l
Desired=Unknown/Install/Remove/Purge/Hold
| Status=Not/Installed/Config-files/Unpacked/Failed-config/\
   Half-installed
|/ Err?=(none)/Hold/Reinst-required/X=both-problems \
   (Status,Err: uppercase=bad)
||/ Name          Version          Description
+++-=============-================-==========================
ii  3dchess       0.8.1-11.1       3D chess for X11
ii  acpi          0.09-1           displays information on ACPI…
ii  acpi-support 0.84              a collection of useful apci events
ii  acpid         1.0.4-1ubuntu11  Utilities for ACPI power mgmnt
ii  adduser       3.80ubuntu2      Add and remove users and groups
[much more output deleted]
```

The first few lines above the actual package listing identify the meaning of the various fields on each line and the symbols that they contain. This information is provided in a fairly cryptic fashion that is mandated by the terminal-oriented output produced by the `dpkg` utility. The entries for package name, package version, and package description are pretty straightforward, but the first field requires some serious explanation. The characters in the first field have the following meaning and possible values:

- The first character in the first field indicates the desired status of the package, which is the state that the packaging system thinks the package should be in. Available indicators in this position are `h` (the package is marked as being on hold, and cannot be updated or removed), `i` (the package should be installed), `p` (the package and all associated configuration information is supposed to be purged), `r` (the package is supposed to be removed, but associated configuration files will be preserved), and `u` (the package has never been installed on this system, so its state is unknown).
- The second character in the first field indicates the actual status of the package on your system. Available indicators in this position are `c` (the configuration files for the package are installed, but the package is not), `f` (the script used to complete the configuration of this package, known as a post-installation script, failed for some reason, and the package is therefore not guaranteed to be correctly installed), `h` (the package is partially installed because the installation process was interrupted), `i` (the package is correctly installed), `n` (the package is not installed), and `u` (the package was retrieved, unpacked, and is partially installed, but its post-installation script was not executed).
- The third character highlights any errors that are associated with the package. Available indicators in this position are a space (no errors, which looks remarkably like there is nothing in this position), `H` (the package has been marked as being on hold by a the package management system itself, which usually means that other packages that this package requires are not installed), `R` (reinstallation is required), and `X` (the package both requires re-installation and has been automatically put on hold by the package management system itself).

Now that I've explained what each character position in the first field of a package entry means, it's easy to see what the ASCII art in the `dpkg` output heading means. This is not exactly the same thing as being "intuitive," but that's hard to do in the limited number of characters that you can display and easily use from the command line.

Listing the Packages That Are Available for Your System

The previous section explained how to figure out what packages are installed on your system. While that's interesting, a more interesting question perhaps is, "What packages are available for my system that I have not yet installed?" As mentioned in the previous section, the `dpkg` command's `-l` option with no other arguments shows the list of installed packages. You can supply an argument to this command to list available packages that match that argument, whether installed or uninstalled. For example, the `dpkg -l emacs` command should list all of the packages that have the string `emacs` in their package name. Let's try that:

```
$ dpkg -l emacs
No packages found matching emacs.
```

Well, that seems odd because I'm actually typing this in emacs on my Ubuntu system. The problem is that any argument that you supply to `dpkg` is used as a pattern match, and there are apparently no packages installed on my system whose exact name is `emacs`. Let's try that again, using a standard Linux wildcard to say that you want to list any packages with a name that begins with `emacs`, as in the following example:

```
$ dpkg -l 'emacs*'
Desired=Unknown/Install/Remove/Purge/Hold
| Status=Not/Installed/Config-files/Unpacked/Failed-config/\
   Half-installed
|/ Err?=(none)/Hold/Reinst-required/X=both-problems \
   (Status,Err: uppercase=bad)
||/ Name                 Version        Description
+++-====================-==============-==========================
ii  emacs-chess          2.0b5-1        a client/library for Chess
ii  emacs-chess-pieces   2.0b5-1        XPM images of chess pieces
ii  emacs-goodies-el     26.4-1         Miscellaneous Emacs add-ons
un  emacs-goodies-extra-el    <none>    (no description available)
ii  emacs21              21.4a-3ubuntu2 The GNU Emacs editor
ii  emacs21-bin-common   21.4a-3ubuntu2 The GNU Emacs shared...
ii  emacs21-common       21.4a-3ubuntu2 The GNU Emacs shared...
un  emacs21-el           <none>         (no description available)
un  emacs21-nox          <none>         (no description available)
un  emacsen              <none>         (no description available)
ii  emacsen-common       1.4.17         Common facilities for…
```

That's more like it! Note that when using wildcards in a `dpkg` specification, you have to protect them from expansion on the command line by enclosing your `dpkg` wildcard specification within quotation marks of some sort.

The previous output showed some Emacs packages, but I could swear that there were others on my Ubuntu system. Let's try a more general wildcard that searches for packages that contain the string `emacs` anywhere in their names:

```
$ dpkg -l '*emacs*'
```

```
Desired=Unknown/Install/Remove/Purge/Hold
| Status=Not/Installed/Config-files/Unpacked/Failed-config/\
  Half-installed
|/ Err?=(none)/Hold/Reinst-required/X=both-problems \
   (Status,Err: uppercase=bad)
||/ Name                 Version         Description
+++-====================-===============-====================
ii  emacs-chess          2.0b5-1         a client/library for Chess
ii  emacs-chess-pieces   2.0b5-1         XPM images of chess pieces
ii  emacs-goodies-el     26.4-1          Miscellaneous add-ons
un  emacs-goodies-extra-el  <none>       (no description)
ii  emacs21              21.4a-3ubuntu2  The GNU Emacs editor
ii  emacs21-bin-common   21.4a-3ubuntu2  The GNU Emacs shared...
ii  emacs21-common       21.4a-3ubuntu2  The GNU Emacs shared...
un  emacs21-el           <none>          (no description available)
un  emacs21-nox          <none>          (no description available)
un  emacsen              <none>          (no description available)
ii  emacsen-common       1.4.17          Common emacsen facilities
un  xemacs               <none>          (no description available)
un  xemacs-support       <none>          (no description available)
un  xemacs-widget        <none>          (no description available)
ii  xemacs21             21.4.18-1ubuntu1 highly customizable editor
[additional output deleted]
```

Even better! Thanks to the explanation of the characters in the first field of your `dpkg` output (provided in the previous section of this chapter), it's clear that I need to install `xemacs` on this system.

Because this is a Linux system, you can pipe your `dpkg` output through other commands to answer questions like, "What packages are not yet installed on my system whose names contain the string `emacs` anywhere in their names," which I can do using `grep` to look for lines that begin with the string `un`, as in the following example:

```
$ dpkg -l '*emacs*' | grep '^un'
un  emacs-goodies-extra-el <none>     (no description available)
un  emacs21-el             <none>     (no description available)
un  emacs21-nox            <none>     (no description available)
un  emacsen                <none>     (no description available)
un  xemacs                 <none>     (no description available)
un  xemacs-support         <none>     (no description available)
un  xemacs-widget          <none>     (no description available)
un  xemacs21-gnome-mule    <none>     (no description available)
un  xemacs21-gnome-mule... <none>     (no description available)
un  xemacs21-gnome-nomule  <none>     (no description available)
un  xemacs21-mule-canna-wnn <none>    (no description available)
un  xemacs21-nomule        <none>     (no description available)
un  xemacs21-supportel     <none>     (no description available)
```

Other Ubuntu package management utilities provide similar search features, particularly the aptitude utility. Although most commonly associated with its terminal-oriented, quasi-graphical interface (discussed later in this chapter in the section "Using aptitude to Add and Remove Software," aptitude also provides a command-line mode with many powerful commands. To locate packages, aptitude provides

a search keyword that enables you to specify a substring to search for (no wildcards necessary), and displays information about any installed or available package that is present in any repository listed in your `/etc/apt/sources.list` file. The following is sample output from a search for Emacs using aptitude:

```
$ aptitude search emacs
p   acl2-emacs          - A Computational Logic for Applicative Comm
p   aleph-emacs         - The Aleph programming language - emacs mod
p   cxref-emacs         - Generates latex and HTML documentation for
p   emacs               - The GNU Emacs editor
i   emacs-chess         - a client and library for playing Chess fro
i   emacs-chess-pieces  - XPM images of chess pieces for emacs-chess
p   emacs-color-themes  - Color themes for Emacs
p   emacs-extra         - emacs configuration
i   emacs-goodies-el    - Miscellaneous add-ons for Emacs
i   emacs21             - The GNU Emacs editor
i   emacs21-bin-common  - The GNU Emacs editor's shared, architectur
i   emacs21-common      - The GNU Emacs editor's shared, architectur
[additional output deleted]
```

You'll notice that the aptitude utility finds more packages than the `dpkg -l` search command that was described earlier.

Another very convenient application for searching for packages is the `apt-cache` utility, which is a command-line utility that provides subcommands that enable you to search and manipulate the cache of packages that are installed or available for installation on your system and the metadata that is associated with those packages. You must enable the universe repository in order to install the `apt-cache` utility. The `apt-cache` utility can also be used to search the packages in any repository that is active in your `/etc/apt/sources.list` file, and find even more matching packages because it searches within the package name and both the short and long descriptions of all available packages, whether installed or not. The following is sample output from an `apt-cache` search for our favorite string, Emacs:

```
$ apt-cache search emacs
acl2-emacs - A Computational Logic for Applicative Common Lisp: emacs
     interface
ada-mode - Ada mode for GNU Emacs and XEmacs
af - An Emacs-like mail reader and composer
aleph-emacs - The Aleph programming language - emacs mode
anjsp - A major mode to edit JSP and PSP code with Emacs
anthy-el - A Japanese input method (elisp fronted)
apel - portable library for emacsen
aplus-fsf-el - XEmacs lisp for A+ development
artist - Emacs Lisp drawing package
asn1-mode - Emacs mode for editing ASN.1 specification files
bhl - Emacs mode for converting brut text to HTML and LaTeX
bigloo-ude - Bigloo Unified Development Environment for Emacs
biomode - [Biology] An Emacs mode to edit genetic data
[additional output deleted]
```

As you can see, the `apt-cache` and `aptitude` utilities are easiest to use when searching for packages, and return the most verbose results.

Listing Information About a Package

Once you've found a package that you're interested in, there are an equivalent number of ways to get more detailed information about the package and its contents. However, because the `apt-cache` and `aptitude` utilities return the best search results, it makes sense to use them to display information about any packages that they've identified. For example, to find out detailed information about the `xemacs21` package that was listed in the previous section, you could use either the `apt-cache show xemacs21` or `aptitude show xemacs21` commands, both of which return almost identical information. The following example shows the output from the `aptitude show xemacs21` command, which is slightly more verbose:

```
$ aptitude show xemacs21
Package: xemacs21
New: yes
State: installed
Automatically installed: no
Version: 21.4.18-1ubuntu1
Priority: optional
Section: universe/editors
Maintainer: OHURA Makoto <ohura@debian.org>
Uncompressed Size: 49.2k
Depends: xemacs21-mule (= 21.4.18-1ubuntu1) | xemacs21-mule-canna-wnn
       (= 21.4.18-1ubuntu1) | xemacs21-nomule (= 21.4.18-1ubuntu1) |
       xemacs21-gnome-mule (= 21.4.18-1ubuntu1) |
       xemacs21-gnome-mule-canna-wnn (= 21.4.18-1ubuntu1) |
       xemacs21-gnome-nomule (= 21.4.18-1ubuntu1)
Conflicts: xemacs, xemacs-widget
Replaces: xemacs, xemacs-widget
Provided by: xemacs21-nomule, xemacs21-mule-canna-wnn,
             xemacs21-mule, xemacs21-gnome-nomule,
             xemacs21-gnome-mule-canna-wnn, xemacs21-gnome-mule
Description: highly customizable text editor
  XEmacs is a full fledged programming language with a mail reader,
news reader, info browser, Web browser, calendar, specialized editor
for more programming languages and other formats than most people
encounter in a lifetime, and much much more.
  This package exists to cause the installation of the real XEmacs
    packages.
```

This output probably provides more than you want to know about the specified package, but it's always better to err on the side of caution.

Listing the Contents of a Package

In some cases, you may want to list all of the files associated with a package that is installed on your system. To do this, you can use the `dpkg` command's `-L` option, followed by the name of the package whose contents you want to list. The following is sample output listing the contents of the `emacs21` package on my Ubuntu system:

```
$ dpkg -L emacs21
/.
/usr
/usr/bin
/usr/bin/emacs21-x
```

```
/usr/share
/usr/share/emacs
/usr/share/emacs/21.4
/usr/share/emacs/21.4/etc
/usr/share/emacs/21.4/etc/DOC-21.4.1
/usr/share/applications
/usr/share/applications/emacs21.desktop
/usr/share/doc
/usr/share/doc/emacs21
/usr/share/doc/emacs21/README.Debian.gz
/usr/share/doc/emacs21/copyright
/usr/share/doc/emacs21/changelog.Debian.gz
/usr/lib
/usr/lib/emacs
/usr/lib/emacs/21.4
/usr/lib/emacs/21.4/i486-linux-gnu
/usr/lib/emacs/21.4/i486-linux-gnu/fns-21.4.1-x.el
/usr/lib/menu
```

Being able to identify the files provided by a package that is installed on your system can be handy if you are considering removing a package but want to make sure that doing so does not delete a file that you want to preserve. The next section describes how to find out what package does provide a specific file.

TIP **You can list the files that are provided by packages that are not yet installed by using the `apt-file` utility with a command such as `apt-file list package`, where `package` is the name of the package that you are interested in. The `apt-file` utility is described in more detail in "Determining What Package Provides a Missing File" later in this chapter.**

Determining What Package Provides an Existing File

After working with Ubuntu for a while and installing, updating, and removing some number of packages, you may be curious about which package provides a certain file on your system. This is easy enough to determine using the `dpkg` command's `--search` option, as in the following example, where I'm curious about which package provides the `/usr/bin/ar` archiving utility:

```
$ dpkg --search /usr/bin/ar
binutils: /usr/bin/ar
```

If you know that a utility that you're looking for is in your path but you don't want to have to determine the directory in which it is physically located, you can use the cool shell trick of using the output of an existing command as the input to another one, in this case combining the `which` command, to determine the full pathname of an application and the `dpkg --search` command, which produces exactly the same output:

```
$ dpkg --search `which ar`
binutils: /usr/bin/ar
```

When searching for files, the `dpkg --search` command can't find symbolic links that are created by package post-installation scripts. If you are trying to find a file that you know exists but cannot locate it, make sure that the file is not a symbolic link and, if it is, search for the file that it actually points to. This can sometimes require a few tries, as in the following example:

```
$ dpkg --search /usr/bin/emacs
dpkg: /usr/bin/emacs not found.
```

```
$ ls -l /usr/bin/emacs
/usr/bin/emacs -> /etc/alternatives/emacs

$ dpkg --search /etc/alternatives/emacs
dpkg: /etc/alternatives/emacs not found.

$  ls -l /etc/alternatives/emacs
/etc/alternatives/emacs -> /usr/bin/emacs21-x

$ dpkg --search /usr/bin/emacs21-x
emacs21: /usr/bin/emacs21-x
```

Although sometimes tedious, symbolic links are usually not an actual part of a package and you must therefore ferret out a file that actually is part of a package.

TIP **Another handy package management utility that you may want to use when determining what package provides a specific file is the `dlocate` utility. This utility is not installed on Ubuntu systems by default, but can easily be installed using `apt-get`, `aptitude`, or Synaptic Package Manager. After installing this command, you can search for the package that provides a file using the `dlocate -S` command, as in the following example:**

```
$ dlocate -S /usr/bin/emacs
emacs21-bin-common: /usr/bin/emacsclient.emacs21
emacs21: /usr/bin/emacs21-x
emacs21: /usr/bin/emacs21
```

The `dlocate` utility's `-S` option uses the string that you specify as a substring to search for, and can therefore be a bit more helpful when looking for files on your system.

Determining What Package Provides a Missing File

Identifying the package that provides a file that is not installed on your system is somewhat tricky, but is a common question when you are trying to build software whose source code requires an include file that is not present on your system, or you are trying to link software that requires a missing library. Unfortunately, the utility that performs this type of search, the `apt-file` utility, is not installed on Ubuntu systems by default, but can easily be installed using `apt-get`, `aptitude`, or Synaptic Package Manager. After installing this command, you must first update its idea of the available packages in the repositories that you are using by running the command `sudo apt-file update`.

Once this command completes, you can search for the package that provides a file, even if that file is not installed on your system, using the `apt-file search` command, as in the following example, which looks for the package associated with the libpowersave library:

```
$ ls -l /usr/lib/*power*
ls: /usr/lib/*power*: No such file or directory
$ apt-file search libpowersave.so
libpowersave-dev: usr/lib/libpowersave.so
libpowersave10: usr/lib/libpowersave.so.10
libpowersave10: usr/lib/libpowersave.so.10.0.3
```

Installing the `apt-file` utility can save you a tremendous amount of hair-pulling and general frustration when you are trying to build software and have no idea what package provides the missing include file or "the missing link." (Sorry, but I couldn't resist.)

TIP **You can also use the `apt-file` utility to list the contents of packages that are not yet installed by using the `apt-file list` *package* command, where package is the name of the package that you are interested in.**

Using apt-get to Add and Remove Software

The apt-get command is the fundamental user-level command provided as part of the Ubuntu package management suite. The apt-get command is a command-line tool that is fast and easy to use.

NOTE **As with any system administrative utility on your Ubuntu system, running the `apt-get` command requires the use of the `sudo` command or an equivalent, such as running `apt-get` under a shell that itself has been executed using the `sudo` command. To simplify examples and the readability of the text, the examples and discussion of using the `apt-get` command throughout this section do not include the `sudo` command.**

Much of the basic functionality provided by the `apt-get` command is quite straightforward. For example, you can install a package and any other packages that it requires using the `apt-get install` package command. Similarly, you can remove a package by using the `apt-get remove package` command, specifying the `--purge` option if you also want to remove any configuration files or other data associated with the package that you're removing. Before using `apt-get` to install new software or new versions of existing software, you should always first execute the `apt-get update` command to ensure that `apt-get` is aware of the latest software packages and versions of that software from all of the repositories listed in your `/etc/apt/sources.list` file.

When using the `apt-get` utility to install software, you may occasionally have to provide additional information during the installation process. In these cases, the `apt-get` utility displays a quasi-graphical screen in the terminal application from which you were running the `apt-get` utility, an example of which is shown in Figure 23-10.

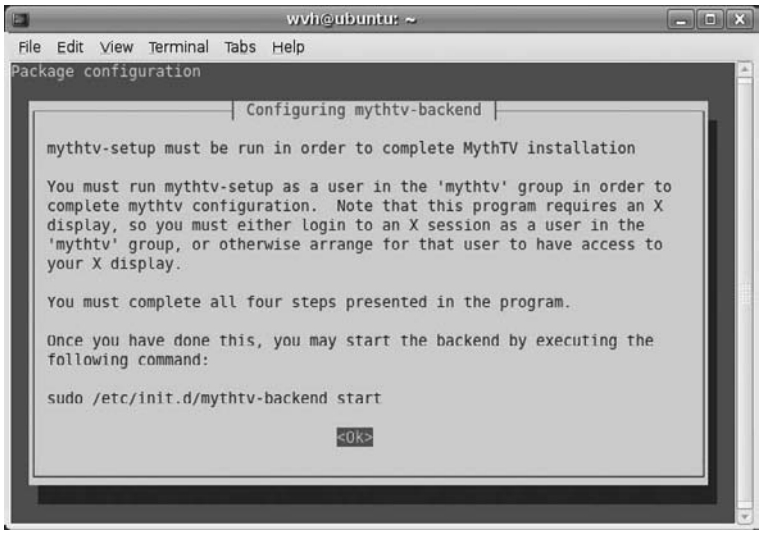
Whenever you see this sort of screen, you must answer any questions that it asks and supply any required information before the installation of your package can complete.

TIP **If you are using `apt-get` to install new software and the `apt-get` utility informs you that other packages are recommended or suggested for installation, you may want to consider using `aptitude`'s command-line interface to install these packages rather than `apt-get`. As explained later in this chapter in the section "Using aptitude to Install Recommended Software" `aptitude` provides options that can automate installing recommended software when installing a new software package.**

As mentioned throughout this book, I prefer to use the graphical Synaptic and Adept Manager tools to search for, install, and remove packages whenever possible. The `apt-get` tool provides these same capabilities, but also provides some powerful capabilities that are not duplicated in the Synaptic and Adept Manager tools, or are at least much more easily done using `apt-get` from the command line. Even if you're a graphical utility fan, the next few sections explain how and when you still may want to use `apt-get` to perform these advanced functions.

FIGURE 23-10

Supplying configuration information during package installation



Upgrading Your System Using apt-get

The `apt-get` upgrade command searches all of the repositories in your `/etc/apt/sources.list` file for new versions of packages that are currently installed on your system, and downloads and installs those new versions. Before using this command, you should always first execute the `apt-get update` command to ensure that `apt-get` is aware of the latest software packages and versions of that software from all of the repositories listed in your `/etc/apt/sources.list` file.

During the upgrade process, the `apt-get upgrade` command will not change the installation status of any other packages on your system. This is the installation status of a package, not the version of a package. If a new version of a package requires a newer version of another package that is already installed on your system, the `apt-get upgrade` command will also install the updated version of that other package. However, if a newer version of an existing package requires that new packages be installed or that existing packages be removed, the `apt-get upgrade` command will not install the newer version of the existing package. To do this, you will need to use the `apt-get dist-upgrade` command, as described in the next section.

Smart System Upgrades Using apt-get

The `apt-get dist-upgrade` command searches all of the repositories in your `/etc/apt/sources.list` file for new versions of packages that are currently installed on your system, and downloads and installs those new versions. Before running this command, you should always first execute the `apt-get update` command to ensure that `apt-get` is aware of the latest software packages and versions of that software from all of the repositories listed in your `/etc/apt/sources.list` file.

The difference between the `apt-get upgrade` and `apt-get dist-upgrade` commands is that the `apt-get dist-upgrade` command will do what is known as a "smart upgrade." This means that it

will do its best to handle any new package requirements or the removal of any existing packages that were mandated in order to install the latest versions of all of the software on your system. For example, if a new version of a package requires that a new package be installed on your system, the `apt-get dist-upgrade` command will install both of these, while the `apt-get upgrade` command would not have updated the original software package. Similarly, if installing a new version of an existing package requires the removal of any other package(s) that are currently installed on your system, the `apt-get dist-upgrade` command will remove those packages in order to install the new version of the existing package.

Even if you are a Synaptic fan, you may occasionally need to execute the `apt-get dist-upgrade` command in order to completely upgrade all of the packages on your system. Synaptic provides a "smart upgrade" preference that tries to do this sort of thing for you whenever possible, but the `apt-get dist-upgrade` command is still occasionally useful. When using Synaptic, you know you need to use the `apt-get dist-upgrade` command when packages are mysteriously identified as being held back and not updated—and, of course, the occasional pop-up message that explicitly tells you to run the `apt-get dist-upgrade` command. The Adept tools refresh your package lists more frequently than Synaptic seems to, and thus are less likely to require using `apt-get` to update things manually—but it's always handy to know how to do so, just in case.

Retrieving Package Source Code Using apt-get

Even in the wonderful world of constantly updated repositories with a rich collection of software, you may occasionally want to build your own versions of software packages for your system. The most common situation in which you'll want to do this is when you find patches to an existing software package that fix problems you're experiencing or add enhancements, and you just can't wait for them to appear in the official or backports repositories. You may also have your own ideas about changes that should be made to a package and want to test and work through your ideas. Linux and its ecosystem of utilities and other applications are open source, after all!

The `apt-get` command makes it remarkably easy to retrieve the source code for an installed package. The only requirement for this capability is that your `/etc/apt/sources.list` file include a `deb-src` entry for the repository from which you retrieved the binary version of that software. For example, if you want to retrieve the source for a package in which the binary version is located in the multiverse repository component, you must have entries like the following (or equivalent entries) in your `/etc/apt/sources.list` file:

```
deb http://archive.ubuntu.com/ubuntu/ intrepid multiverse
deb-src http://archive.ubuntu.com/ubuntu/ intrepid multiverse
```

Assuming that the right entries are present in your `/etc/apt/sources.list` file, retrieving the source code for a specified package is easy using the `apt-get` source command. The following example shows the retrieval of the source code for one of my favorite packages, the MythTV personal video recorder package:

```
$ apt-get source mythtv
Reading package lists... Done
Building dependency tree... Done
Need to get 9836kB of source archives.
Get: 1 http://.../multiverse mythtv 0.21.0+fixes18722-0ubuntu1...
Get: 2 http://... intrepid/multiverse mythtv 0.21.0. . .
Get: 3 http://... intrepid/multiverse mythtv 0.21.0. . .
Fetched 9836kB in 1m31s (107kB/s)
dpkg-source: extracting mythtv in mythtv-0.21.0
dpkg-source: unpacking mythtv_0.21.0.orig.tar.gz
dpkg-source: applying ./mythtv_0.21.0+fixes18722-0ubuntu1.diff.gz
```

When retrieving the source code for a package, the `apt-get source` utility actually retrieves multiple files. The extensions of these files are as follows:

- `diff.gz`: A compressed file containing any available patches to the standard source archive for the package
- `DSC`: A description file for the Ubuntu package that identifies the contents of the package, associated binary packages, any package dependencies for that package, and so on
- `tar.gz`: A compressed archive file containing the source code for the official version of this package

After retrieving these files, the `apt-get source` command unpacks the compressed source archive, creating a working directory for your development efforts, and automatically applies any patches that are available in the (optional) compressed patch file.

Bonus Chapter 3 discusses how to install the basic packages required for software development on an Ubuntu system. If you are writing your own code or simply compiling existing applications on an Ubuntu system, that may be sufficient—except when you are building complex graphical software or when you are building your own DEB packages from the downloaded source for an Ubuntu package. In these cases, you will quickly encounter frustration such as the fact that today's graphical software has dependencies on libraries, include files, and utilities that you may not be familiar with or even have heard of before. Similarly, building DEB packages requires the use of billions and billions of utilities that may be new to you. Not to worry—you're not the first person to have encountered these issues, and the `apt-get` command provides a handy solution, as explained in the next section.

Satisfying Build Dependencies Using apt-get

Many of the software packages available for Linux have dependencies on other packages, which makes perfect sense—why reinvent the wheel when you can just link your code with it? However, if you decide to work on an existing Linux package, build-time dependencies can be frustrating. They're rarely well-documented, so you tend to encounter them as an iterative set of errors when you try to compile and link the application you're trying to build. This is also true when you begin to work with software packages in the formats used by different package management systems. The DEB packages used on Ubuntu and Debian systems (and derivatives) eliminate the dependency mumbo-jumbo that you often encounter with other package formats (such as RPM), but require the presence of an entire ecosystem of related utilities in order to build and package them successfully. Getting started with these can be tricky—you don't know what you're missing until your build process goes up in flames.

The `apt-get` tool provides a great, automated solution for these sorts of problems through its `build-dep` command-line option. Using this command-line option, followed by the name of the package that you want to build automatically identifies, retrieves, and installs all of the include files, libraries, and tools required to build the specified package. The following is an example of the output from this command when retrieving build dependencies for the MythTV package used as an example in the previous section:

```
$ sudo apt-get build-dep mythtv
Reading package lists... Done
Building dependency tree... Done
The following NEW packages will be installed
  build-essential g++ g++-3.4 g++-4.0 libartsc0-dev libasound2-dev
  libaudio-dev libdvb-dev libgl1-mesa-dev libglu1-mesa-dev
  libjpeg62-dev liblame-dev liblcms1-dev liblircclient-dev
  libmng-dev libmysqlclient14-dev libogg-dev libqt3-headers
  libqt3-mt-dev libstdc++6-4.0-dev libstdc++6-dev libvorbis-dev
```

```
    libxmu-dev libxmu-headers libxt-dev libxv-dev libxvmc-dev
    libxxf86vm-dev mesa-common-dev qt3-dev-tools x11proto-video-dev
    x11proto-xf86vidmode-dev
0 upgraded, 32 newly installed, 0 to remove and 0 not upgraded.
Need to get 14.2MB/17.9MB of archives.
After unpacking 63.4MB of additional disk space will be used.
Do you want to continue [Y/n]? Y
[much output deleted]
```

As you can see from this example, it would have taken me many, many iterative compilation attempts to retrieve all of the packages required to build MythTV. Identifying the packages that provide missing include files and libraries can also be time-consuming, even with the help of the `apt-file` application discussed earlier in this chapter in the section "Determining What Package Provides a Missing File." The `apt-get build-dep` command eliminates these sorts of hassles, letting you focus on the software, not the infrastructure required for building it.

Using aptitude to Add and Remove Software

As mentioned earlier in this chapter, the `dpkg` utility provides the conceptual underpinnings of much of the package management software provided on Ubuntu (and Debian) systems. However powerful it may be, it is a command-line utility, and therefore doesn't provide any sort of graphical interface, so the `dselect` utility was developed in order to at least provide a quasi-graphical interface using terminal-oriented cursor-movement libraries such as `curses` (now `ncurses`, for new curses). However, `dselect` couldn't quite let go of its command-line roots, and therefore also provides a command-line interface.

The `apt-get` application discussed earlier in this chapter provides a simpler, higher-level interface than that provided by the `dpkg` utility. However, `apt-get` is YACLU (Yet Another Command-Line utility), and therefore doesn't provide any sort of graphical interface either. The `aptitude` application was developed to solve this problem, but couldn't quite let go of the command line either. Thus `aptitude` (like `deselect` before it) provides both a powerful command-line interface and a terminal-oriented, quasi-graphical interface.

Personally, I find terminal-oriented graphical interfaces quaint, at best. I prefer not to use them unless absolutely necessary, especially when they've been replaced by actual graphical user interfaces such as that provided by the Synaptic and Adept Package Managers.

NOTE **When using the `aptitude` utility to install software, you may occasionally have to provide additional, configuration-related information during the installation process. In these cases, the `aptitude` utility displays the same sort of information requests as those displayed by the `apt-get` utility and shown earlier in Figure 23-10. Whenever you see this sort of screen, you must answer any questions that it asks and supply any requested information before the installation of your package can complete.**

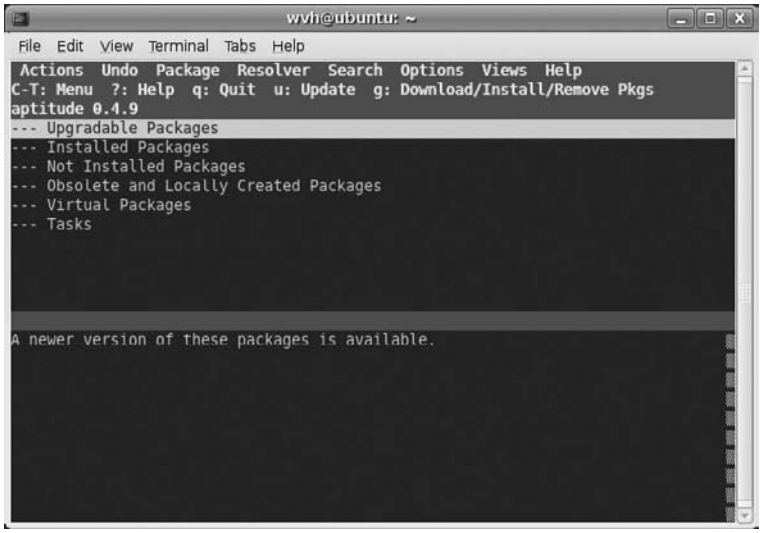
Software can easily be installed by `aptitude` using the command-line `aptitude install` command, and removed using the `aptitude remove` command. As discussed later in the section entitled "Using aptitude to Install Recommended Software," `aptitude` provides some convenient options that make it attractive to use `aptitude` from the command line in certain cases. Many people are also fans of the `aptitude` interface, primarily those who understand it. The next section discusses the absolute basics of the `aptitude` user interface, focusing on how to make it readable by mere mortals.

Tips and Tricks for Using the aptitude User Interface

Figure 23-11 shows the default `aptitude` interface that you see the first time that you start `aptitude` with no arguments in a GNOME Terminal.

FIGURE 23-11

The default aptitude interface in the GNOME Terminal



This interface looks essentially the same in an X Window system xterm or any other terminal application that supports color. (It's slightly more psychedelic in an xterm, but I won't bore you with a screenshot.)

If you insist on using this interface but are having problems seeing the text or dealing with the color scheme, you're somewhat out of luck because the `aptitude` application doesn't offer command-line options for changing the color scheme. However, all is not lost—you can make the interface easier to read by changing your terminal application's underlying notion of the type of terminal in which you are running `aptitude`. A good choice for this is the value `vt100`, which is a classic terminal from the late Digital Equipment Corporation that only supported black and white. To make this change, exit `aptitude` by pressing Q and using the tab key to emphasize that you do actually want to exit from `aptitude`. Next, type the following command in your terminal application:

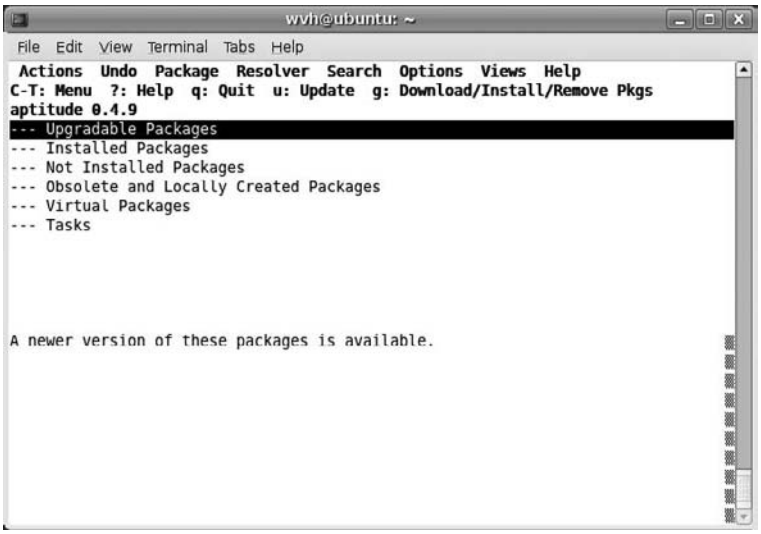
```
$ export TERM=vt100
```

After executing this command, you can then restart `aptitude`, which will look like the screen shown in Figure 23-12.

Changing the color scheme doesn't make the `aptitude` interface any more attractive, but (for me) it does at least make it more usable.

FIGURE 23-12

The aptitude interface in black and white



Once `aptitude` is running, the following are my favorite commands:

- `Q`: Enables you to exit `aptitude`.
- `?`: Displays a panel containing summary information about using the `aptitude` interface.
- `Ctrl+T`: Holding down the Ctrl key and pressing T displays `aptitude`'s menus, which you can navigate through using the arrow keys or by pressing the highlighted key shortcut to execute a specific command.

Figure 23-13 displays a menu in `aptitude` using `aptitude`'s default color scheme. Need I say more? This is the sort of thing that gives Linux applications a bad name to people who are used to Macintosh (or even Windows) graphics.

For more detailed (and less opinionated) information about using the `aptitude` user interface, see the online reference information for `aptitude` (`man aptitude`) or search the Web for relevant information.

Using aptitude to Install Recommended Software

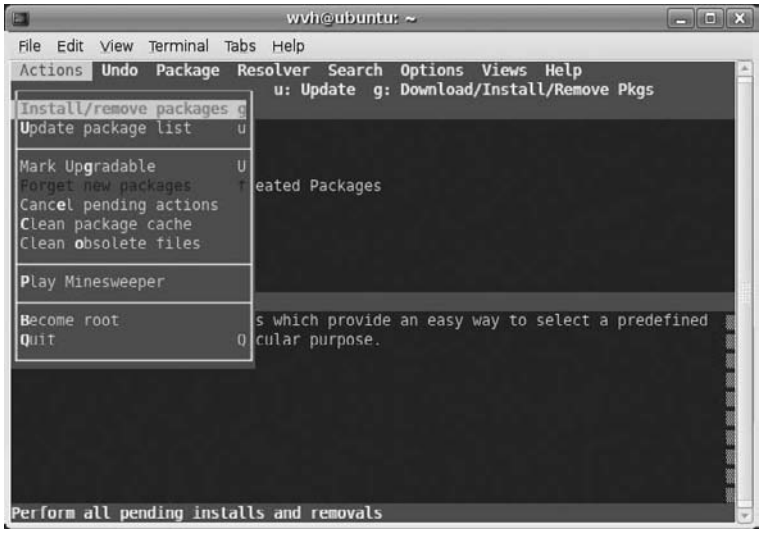
If you have been using `apt-get` to install new software and the `apt-get` utility informs you that other packages are recommended or suggested for installation, you may want to consider using `aptitude`'s command-line interface rather than `apt-get`. The DEB packages used by Ubuntu identify various dependences, divided into five different classes:

- **Conflicts:** Packages (or package versions) with which a package conflicts, and which cannot therefore be installed on a system when a specified package is being installed.

- **Depends:** Packages that are mandatory for the correct operation of a package, and which therefore must also be installed when a given package is installed.
- **Enhances:** Packages whose operation is improved or simplified as a result of installing a package, but between which there is no direct relationship.
- **Pre-Depends:** Packages that must be completely installed on your system before you can install a given package.
- **Recommends:** Packages that you should probably have, but which are not absolutely required for the correct operation of a given package.
- **Suggests:** Packages that may be useful, and which are normally installed on systems where a given package is installed.

FIGURE 23-13

A menu in the default aptitude interface



Requirements marked as Pre-Depends (which are rarely used or necessary) must be satisfied before you can install a given package. Requirements marked as Depends are automatically installed along with a given package. The `aptitude` utility enables you to automatically install packages that are recommended for installation by including the `--with-recommends` command-line option on your `aptitude` install command line.

The aptitude utility is the only command-line package installation utility that enables you to automatically install recommended packages along with required packages. Specifying the `--with-recommends` command-line option when installing packages from the command line can be quite convenient and can save you search time in the future if you discover that you want to use a recommended package. A similar option to aptitude's `--with-recommends` option for `apt-get`, `--recommends`, and a related on for

installing suggested packages, `--suggests`, have been "suggested" for `apt-get` for a long time, but are not available in any version of `apt-get` that I've ever seen. As I'll discuss in the section of this chapter on Synaptic, the Synaptic tool provides configuration options for considering recommended packages to be dependencies and automatically installing them; similarly, its interface provides a run-time option for manually selecting suggested packages and including them as part of the package installation process.

Advantages of Using aptitude to Install and Remove Software

Although I am not such a fan of the `aptitude` user interface, `aptitude` itself does provide some significantly useful capabilities. Its search capabilities, discussed earlier in this chapter in the section "Listing the Packages That Are Available for Your System," are easy to use and quite powerful—exactly what you want in a utility. Similarly, the current version of `aptitude` has some other advantages when installing and, specifically, when removing packages. The most significant of these is the fact that `aptitude` remembers dependency information when installing packages. Every Ubuntu package management utility understands requirements when installing packages, but `aptitude` remembers this information and can therefore use that information should you decide to remove packages.

The following sample output shows an attempt to remove the PostgreSQL database system from one of my Ubuntu systems using `apt-get`:

```
$ sudo apt-get remove postgresql
Password:
Reading package lists... Done
Building dependency tree... Done
The following packages will be REMOVED
  postgresql
0 upgraded, 0 newly installed, 1 to remove and 0 not upgraded.
Need to get 0B of archives.
After unpacking 45.1kB disk space will be freed.
Do you want to continue [Y/n]? n
Abort.
```

Using `aptitude` to remove the same package is much more useful because it also offers to remove packages that were installed along with the `postgresql` package, but which are no longer needed on that system if the `postgresql` package itself is removed, as in the following example:

```
$ sudo aptitude remove postgresql
Reading package lists... Done
Building dependency tree... Done
Reading extended state information
Initializing package states... Done
Building tag database... Done
The following packages are unused and will be REMOVED:
  postgresql-plperl-7.4 postgresql-plpython-7.4 postgresql-pltcl-7.4
The following packages will be REMOVED:
  postgresql
0 packages upgraded, 0 newly installed, 4 to remove \
  and 0 not upgraded.
Need to get 0B of archives. After unpacking 651kB will be freed.
Do you want to continue? [Y/n/?] Y
[much output deleted]
```

For me, this makes `aptitude` a much more attractive command-line package installation solution than `apt-get`. Note that `aptitude` only remembers dependency information for packages that it has installed, so I tend to use `aptitude` from the command line for both package installation and removal. Otherwise, if you actively install and remove large numbers of packages, your system tends to accumulate packages that you no longer need, but which you're not aware of. For information about ways of checking for unneeded packages and removing them, see the last section of this chapter, "Keeping your System Lean, Mean, and Pristine."

Using Synaptic to Add and Remove Software

The Synaptic Package Manager is a graphical, GNOME application for installing, removing, and generally managing Ubuntu software packages. In the unlikely event that this is only the second sentence you have read in this book, you may not know that I recommend using the Synaptic package management utility for almost all package management tasks on Ubuntu systems. It's true. Synaptic is the type of administrative application that users of other Linux distributions and other personal computer operating systems wish they had.

Synaptic was originally developed by the folks at Connectiva Linux, a Latin American Linux distribution that was eventually acquired by Mandrake Linux, which then changed its name to Mandriva to highlight that it wasn't your father's Mandrake Linux distribution anymore. Synaptic was originally written using the WINGS toolkit used by the Window Maker X Window system window manager (`www.windowmaker.info`), which attempted to provide a look and feel familiar to users of NeXT computer systems. Today, Synaptic is built with the standard GTK+ toolkit used by GNOME. Interestingly, Connectiva was an RPM-based distribution, but Synaptic is now more commonly associated with the DEB packages used by distributions such as Ubuntu and Debian.

The Synaptic package management utility is installed by default as part of any Ubuntu installation. Select the System ➪ Administration ➪ Synaptic Package Manager menu item to start the Synaptic tool. As when trying to run any graphical administrative application on Ubuntu, you will be prompted for your administrative password before Synaptic actually starts. Figure 23-14 shows the dialog displayed when you first start Synaptic.

The left pane of the Synaptic dialog shown in Figure 23-14 shows the conceptual categories into which software is organized based on the default repositories enabled when you first install Ubuntu.

The basic package-related areas of Synaptic's main dialog are as follows:

- **Left pane:** Organizes available software into various conceptual categories based on the repositories listed in `/etc/apt/sources.list`.
- **Upper-right pane:** Lists the software packages that are available in the category that is currently selected in the left pane. By default, the All category is selected in Synaptic's left pane, so that all available software packages are listed. As I discuss later in this chapter, performing a search in Synaptic creates a new category with the name of your search term(s) and hides all other default categories with the exception of the All category.
- **Bottom-right pane:** Displays the description of any package that is currently selected in the upper-right pane.

Synaptic provides an excellent online manual, shown in Figure 23-15, that you can access by selecting the Help menu's Contents item.

FIGURE 23-14

The startup dialog for the Synaptic utility

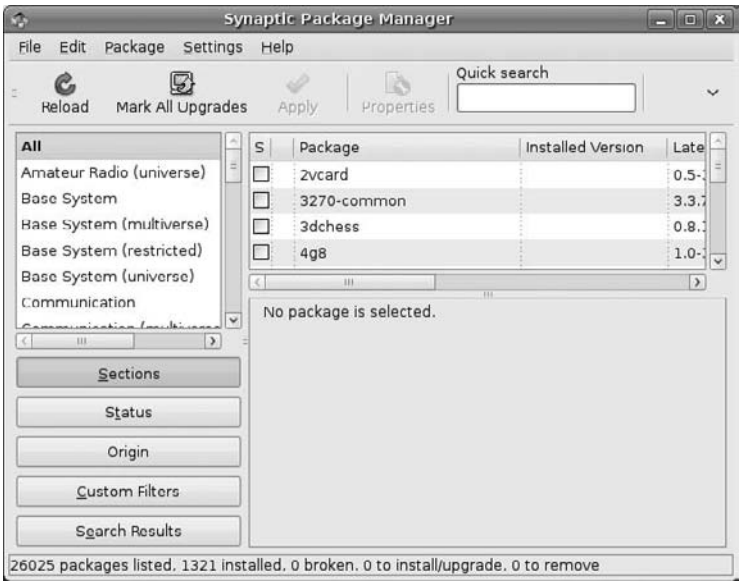
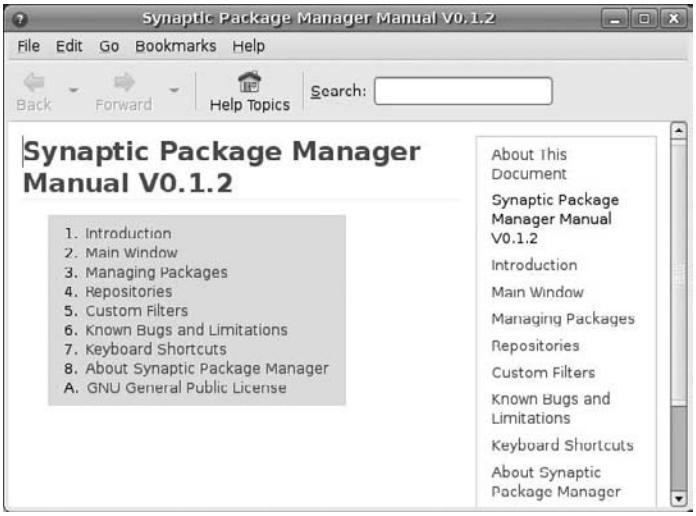


FIGURE 23-15

The online manual for Synaptic



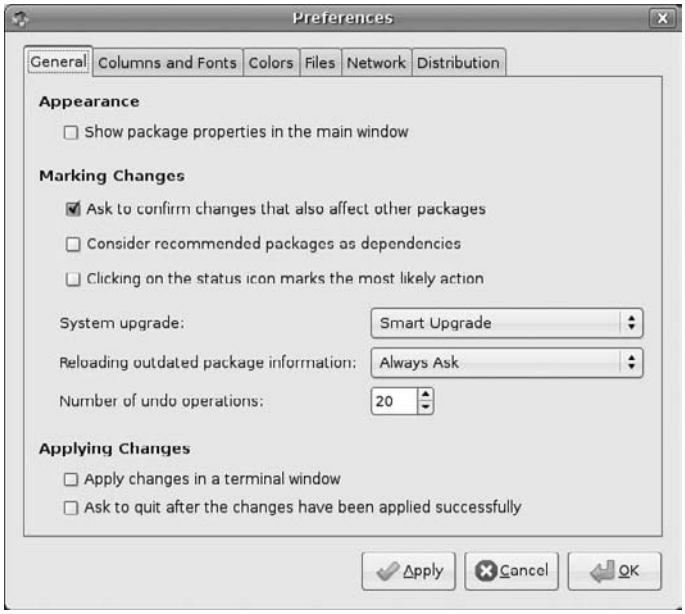
Because Synaptic includes thorough and up-to-date online help, I'm not going to bore you by repeating all of that. The next few sections focus on how to do common tasks in Synaptic, explaining how to configure Synaptic and how to perform basic package management tasks such as searching for and installing packages using Synaptic.

Configuring Synaptic Preferences

The entries on Synaptic's Settings menu enable you to set default values for using Synaptic, identify the repositories available to Synaptic, configure how searches work in Synaptic, configure how the icons are displayed in the Synaptic toolbar, and set internal variables used by Synaptic. Most of the settings that you will want to change are located on the Preferences dialog, which you can display by selecting the Settings menu's Preferences option. Figure 23-16 shows Synaptic's Preferences dialog.

FIGURE 23-16

The Synaptic Preferences dialog



The Synaptic preferences that you will most commonly want to change are all located on the general tab of Synaptic's Preferences dialog. These are the following:

- **Consider recommended packages as dependencies:** Selecting this checkbox causes Synaptic to always install any packages that are associated with a package that you are installing, but which are identified as recommended, but not required, for package installation. As discussed in more detail in this chapter in the section "Using aptitude to Install Recommended Software," such recommended packages are ones that you will probably want to use with the package that you are installing.

- **System upgrade:** This item should always be set to "Smart Upgrade," which tries to identify other packages that must be installed, updated, or removed when installing or upgrading any package. Synaptic is much more conservative about this than other utilities such as `apt-get` (where the `apt-get dist-upgrade` command is a more powerful equivalent for this Synaptic option. Even if you are a Synaptic fan, you will occasionally need to execute the `apt-get dist-upgrade` command in order to completely upgrade all of the packages on your system. When using Synaptic, you know you need to use the `apt-get dist-upgrade` command when packages are mysteriously identified as being held back and not updated—and, of course, the occasional pop-up message that explicitly tells you to run the `apt-get dist-upgrade` command.
- **Apply changes in a terminal window:** Selecting this option causes Synaptic to always display package installation/update status messages and results in a separate terminal window. As mentioned when discussing installing software using the `apt-get` and `aptitude` utilities, you may occasionally have to provide additional, configuration-related information during the installation process. In these cases, the Synaptic utility displays the same sort of information requests as those displayed by the `apt-get` and `aptitude` utilities, as shown earlier in Figure 23-10. Whenever you see this sort of screen, you must answer any questions that it asks and supply any requested information before the installation of your package can complete. The Synaptic utility displays this information in the dialog that you see only when you expand the software installation status window, so it's a good idea to always monitor this window. If you choose not to and notice that Synaptic seems to have stopped installing things, expand this window to see if Synaptic is waiting for additional information.

Once you have changed any setting on any of the tabs in Synaptic's Preferences dialog, click Apply to make sure that your changes are saved, and then click OK to close the Preferences dialog and return to the main Synaptic dialog.

Searching for Software in Synaptic

One of Synaptic's most powerful capabilities is providing a powerful and usable package search capability that doesn't involve a command line. To locate packages that are relevant to one or more keywords, you can either use the Quick Search text box in the main Synaptic dialog to search package descriptions and names, or click the down arrow to the right of the Quick Search area and select the Search button to display a more advanced search dialog that also enables you to specify what package fields to search, as shown in Figure 23-17.

FIGURE 23-17

Synaptic's Find dialog



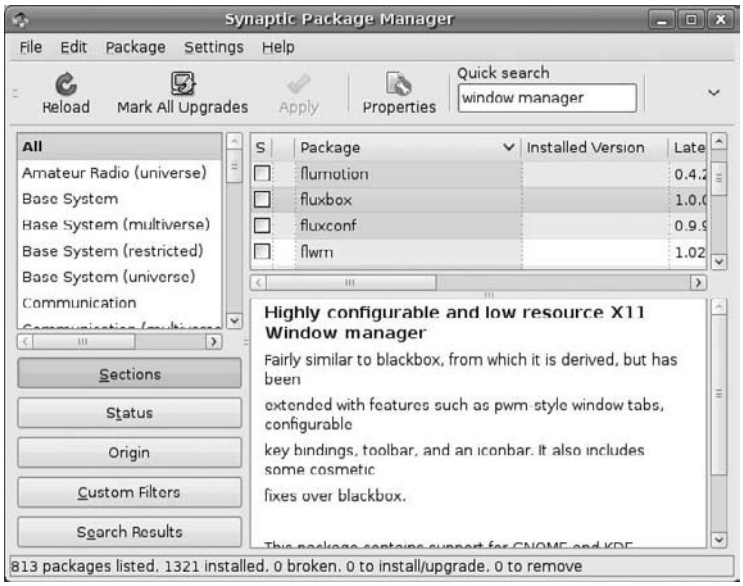
Most examples of installing software throughout this book have relied on Synaptic's search capabilities. To locate matching, available packages, you have to make sure that you're searching the right portion of the information that Synaptic has about packages. In almost all cases, you will want to make sure that the Look in value in Synaptic's Search dialog is set to Description and Name so that the search looks for matches for your keyword(s) in both package descriptions and package names. However, you may occasionally want to search other portions of the package information that is available via Synaptic. Other search possibilities are as follows:

- **Name:** Searches only package names for your keyword
- **Maintainer:** Looks for a package maintained by one your favorite Ubuntu rock stars
- **Version:** Finds specific versions, which can be useful when looking for all packages that take their version number from higher-level packages, such as GNOME versions
- **Dependencies:** Identifies packages that depend on a given package
- **Provided Packages:** Identifies packages based on other packages that they also install

To search for packages in Synaptic, enter one or more keywords in Synaptic's Quick Search text area or Find dialog and click Search (or press Return). Synaptic will create a new category in the left pane with the name of your search term(s) and hide all other default categories with the exception of the All category. Once the search completes, the new category is selected and any matching packages display in the upper-right pane. Figure 23-18 shows the results of a search for "window manager."

FIGURE 23-18

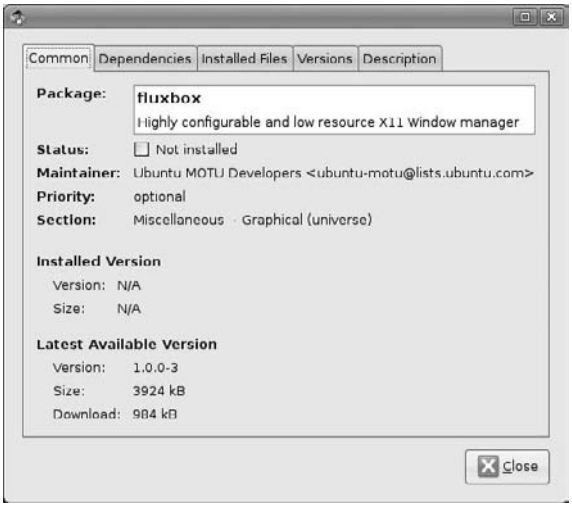
Search results in Synaptic



To select a package for installation, scroll down in the upper-right pane until you see the name of the package that you want to install. You can select the name of any package in the upper-right pane to see the description of that package, and can also click the Properties button in Synaptic's toolbar (or select the properties item from the pop-up menu after right-clicking its name) to display a Properties dialog that shows detailed information about the selected item. Figure 23-19 shows the Properties dialog for the package that provides the fluxbox window manager, a lightweight window manager that I like to use on laptops.

FIGURE 23-19

A package Properties dialog in Synaptic



The Common tab, shown by default when you display a package properties dialog, provides general information about the package. The Dependencies tab provides useful information about required, recommended, suggested, and conflicting packages. The Installed Files tab is also very useful, but only contains information on packages that are already installed. (To see information about the files that are included in a package that is not installed, you can use the `apt-file` application, which is discussed earlier in this chapter in the section "Determining What Package Provides a Missing File.")

After examining any properties of a selected package, click Close to close the properties dialog and return to the main Synaptic dialog.

Installing Packages in Synaptic

After locating a package that you want to install, right-click on its name in the upper-right pane and select Mark for Installation from the pop-up menu, as shown in Figure 23-20.

If installing the selected package requires that other packages be installed, Synaptic displays a dialog like that shown in Figure 23-21.

FIGURE 23-20

Marking a package for installation in Synaptic



This dialog lists all packages whose installation is required by the package that you have selected. If you checked the Consider recommended packages as dependencies item in Synaptic's Preferences dialog, this list will also include any packages that are recommended, but not required, for installation. Click Mark to accept the installation of these other packages and close this dialog.

Next, right-click the package name again and check whether any packages are listed on a submenu of the Mark Suggested for Installation command at the bottom of Synaptic's pop-up menu, as shown in Figure 23-22.

FIGURE 23-21

Identifying other packages required for installation

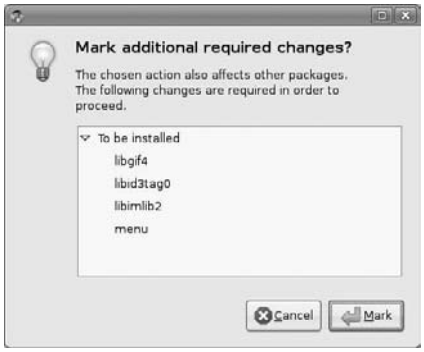
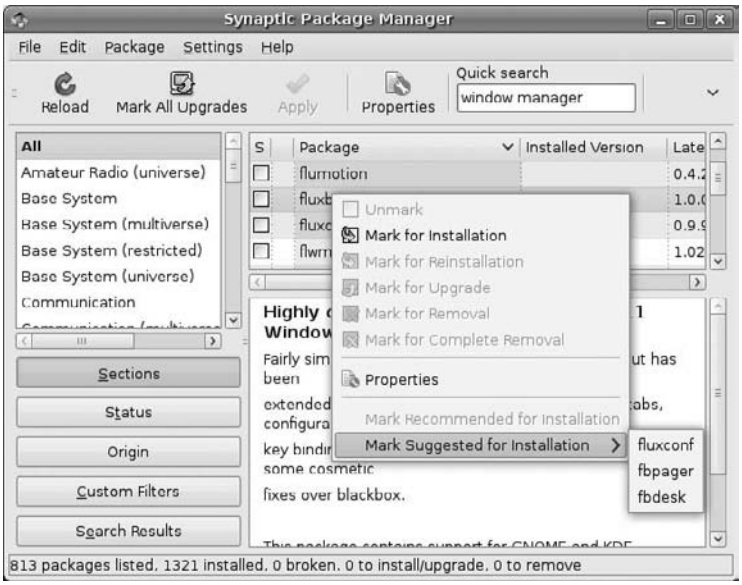


FIGURE 23-22

Marking suggested packages for installation in Synaptic



This submenu identifies packages that are normally installed with the primary package that you have selected for installation, but which are not required or expected. You can select any of the additional packages listed on this menu (one by one) to also schedule them for installation.

Once you have selected any suggested packages that you want to install with the primary package that you have selected for installation, click Apply in the Synaptic toolbar to begin the actual installation process. This displays the dialog shown in Figure 23-23, which summarizes all of the activities associated with installing the selected package(s).

To proceed, click Apply. At this point, Synaptic begins retrieving the selected packages from the repository where they are located, displaying a progress dialog like the one shown in Figure 23-24.

Once the download process completes, Synaptic begins installing and configuring the downloaded packages. If you did not select the Apply changes in a terminal window item in Synaptic's Preferences dialog, Synaptic displays a simple status dialog, as shown in Figure 23-25.

You can display a small terminal window that lists the status of Synaptic's installation and configuration processes by clicking the arrow to the right of the Terminal option on this dialog. When the installation and configuration of your packages completes, the dialog shown in Figure 23-26 displays.

FIGURE 23-23

A package installation summary in Synaptic

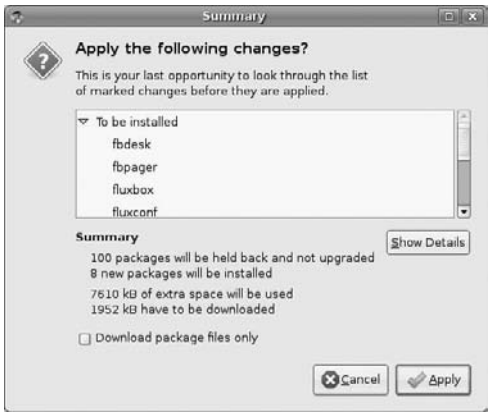


FIGURE 23-24

The download status dialog during Synaptic package installation



TIP **When using the Synaptic utility to install software without having configured Synaptic to always display configuration status in a Terminal window, it is useful to display a Terminal to see information about the status of software installation on your system. You may occasionally have to provide additional, configuration-related information during the installation process. In this case, the Synaptic utility displays the same sort of information requests as those displayed by the `apt-get` and `aptitude` utilities, as shown earlier in Figure 23-10. Whenever you see this sort of screen, you must answer any questions that it asks and supply any requested information before the installation of your package can complete. The Synaptic utility displays this information in the dialog that you see only when you expand the software installation Terminal window, so it's a good idea to always monitor this window. If you choose not to and notice that Synaptic seems to have stopped installing things, expand this window to see if Synaptic is waiting for additional information.**

FIGURE 23-25

A simple installation and configuration status dialog in Synaptic



FIGURE 23-26

A simple completion dialog in Synaptic



If you selected the Apply changes in a terminal window item in Synaptic's Preferences dialog, Synaptic displays a single large terminal window during its installation and configuration phases, shown in Figure 23-27, instead of the dialogs shown in Figures 23-25 and 23-26.

Regardless of which of the dialogs you see when package installation and configuration completes, click Close to close that dialog and return to Synaptic's main dialog.

Once you have finished installing any packages that you are interested in, you can exit from Synaptic by selecting the File menu's Quit command (or pressing the Ctrl+Q key sequence).

Removing Packages in Synaptic

Removing software in Synaptic is very similar to the installation process. You must first locate the package(s) that you want to remove by selecting the appropriate category from Synaptic's left pane and scrolling through the package lists in the upper-right pane, or by searching for that package and then selecting it from the upper-right pane. Either way, once you have located the package that you want to remove, right-click on its name to display the pop-up menu shown in Figure 23-28.

Marking a package for complete removal removes the package and any configuration files for the application that you may have modified, while Mark for Removal removes the package but leaves the configuration files on your system.

FIGURE 23-27

A terminal window showing installation and configuration status

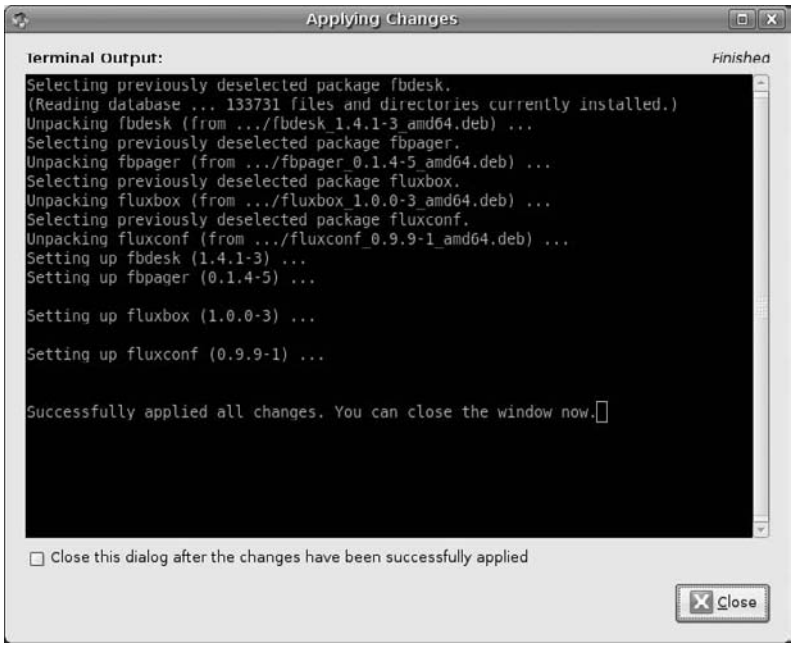


FIGURE 23-28

Marking a package for complete removal in Synaptic

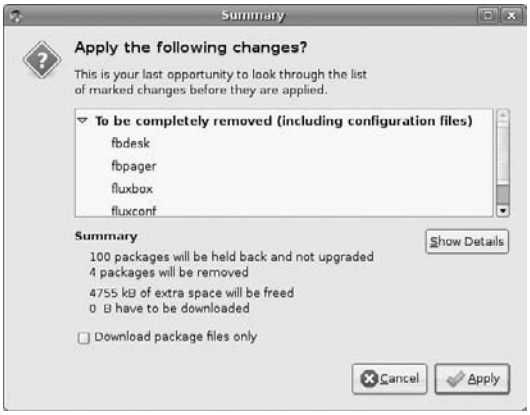


TIP **Unfortunately, Synaptic does not automatically remove other packages that were required by a package that you are selecting for removal, or which you installed as a result of their being suggested for installation. You must select all of these packages manually in order to remove them. You can also not worry about them and use the process described in the last section of this chapter, "Keeping your System Lean, Mean, and Pristine," to remove them at some point in the future.**

Once you have selected any packages that you want to remove and marked them for removal, click Apply to proceed with the installation process. This displays the confirmation dialog shown in Figure 23-29.

FIGURE 23-29

The confirmation dialog for package removal in Synaptic



Click Apply to continue. As the removal process proceeds, the Synaptic dialog displays status dialogs similar to those shown during the package installation process, as described in the previous section. Once package removal completes, click Close to close the final status dialog and return to Synaptic's primary dialog.

Once you have finished removing packages from your system, you can exit from Synaptic by selecting the File menu's Quit command (or pressing the Ctrl+Q key sequence).

Using the Ubuntu Update Manager

Ubuntu's Update Manager checks the active repositories listed in your `/etc/apt/sources.list` file and notifies you if updates to any installed packages are available. The Update Manager uses an associated panel applet to let you know when updates are available by displaying a small orange icon and a pop-up dialog when updates are first found. If you see this icon, you can move the cursor over it at any time to display a tooltip that summarizes available updates. Figure 23-30 shows this icon and the tooltip that provides information about the availability of system updates.

Clicking on this icon starts the Update Manager itself, which is shown in Figure 23-31.

FIGURE 23-30

The icon and dialog that notify you of available updates

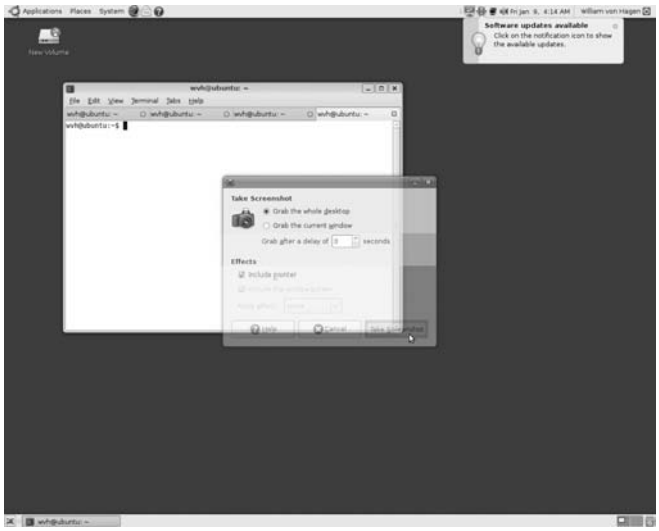
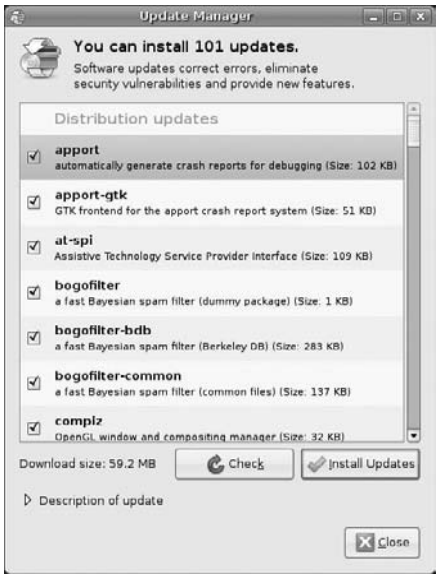


FIGURE 23-31

The Ubuntu Update Manager



Using the Update Manager is simple—just click Install Updates. The Update Manager then automatically retrieves, installs and configures all available updated packages, displaying progress and status dialogs that are very similar to those used by the Synaptic Package Manager. When the update process completes, click Close to terminate the Update Manager and begin doing real work.

TIP **As when using the Synaptic utility, you should watch for configuration questions or requests for additional information that are required during the package update process. If the update process seems to have halted or is taking a long time, click the Terminal button in the Update Manager's status dialogs to see if any configuration questions are pending.**

Using the Adept Tools

The Adept Package Manager is a graphical KDE application for installing, removing, and generally managing Kubuntu software packages. Like its GNOME-based equivalent, Synaptic, the Adept Package Manager and its friends, the Adept Installer and the Adept Updater, are the types of administrative application that users of other Linux distributions and other personal computer operating systems wish they had.

The lead developer for the Adept utilities is Peter Rôckai, and their development is being sponsored by Canonical. This is a really Good Thing™ because all of the graphical installers for KDE that I'm aware of are RPM-oriented, and thus don't work well with the Debian package format used by Kubuntu. A KDE-based version of Synaptic called Kynaptic was under development, but seems to have died off. Thankfully, Canonical has stepped up to the plate to fund the Adept utilities and provide a set of DEB-friendly graphical package management utilities for Kubuntu that can also be used by any other KDE-based Linux distributions that use the Debian package format. Yet another reason to admire Canonical and their contributions to computing today.

The versions of the Adept utilities found on KDE 3-based and KDE 4-based systems look very different and work very different. KDE 4 is providing an opportunity to not only rethink the appearance of these tools, but to also rethink their operation. The next two sections discuss how to use the KDE 3 and KDE 4 versions of the Adept Package Manager and Adept Updater utilities, so that you can install all of the cool Kubuntu software that you need, and keep that software (and your system itself) up-to-date as updates and enhancements to Kubuntu appear, regardless of which version of Kubuntu you're using.

Using the Adept Manager on KDE 3-Based Systems

The Adept utilities are installed by default as part of any Kubuntu installation. Select the K Menu ➪ System ➪ Adept Manager menu item to start the Adept Package Manager. As when trying to run any graphical administrative application on Kubuntu, you will be prompted for your administrative password before the Adept Manager actually starts. Figure 23-32 shows the dialog displayed when you first start the Adept Manager.

From top to bottom and left to right, the different areas of the main Adept Manager dialog are as follows:

- **Menus and toolbar:** Provide easy access to specific Adept commands, configuration mechanisms, and online help.
- **Active Filters pane:** Displays the filters that you can use to locate a specific package or refine the set of available packages that is currently being displayed.
- **Package List pane:** Lists the software packages that are available based on the filters that are currently being used. To display a more verbose description of any package listed in this pane, click the blue arrowhead to the left of the name of that package, as shown in Figure 23-33. To close the verbose description, click the arrowhead again.

FIGURE 23-32

The startup dialog for the KDE 3 Adept Manager

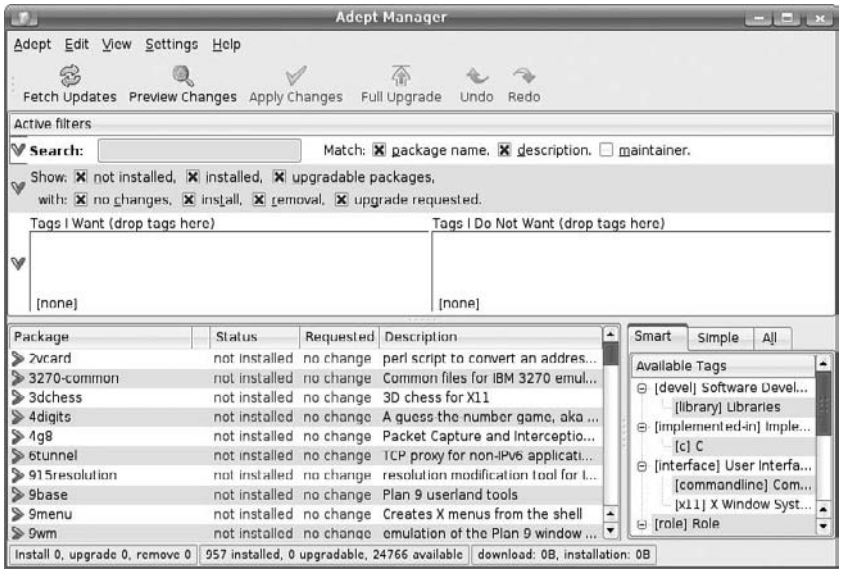
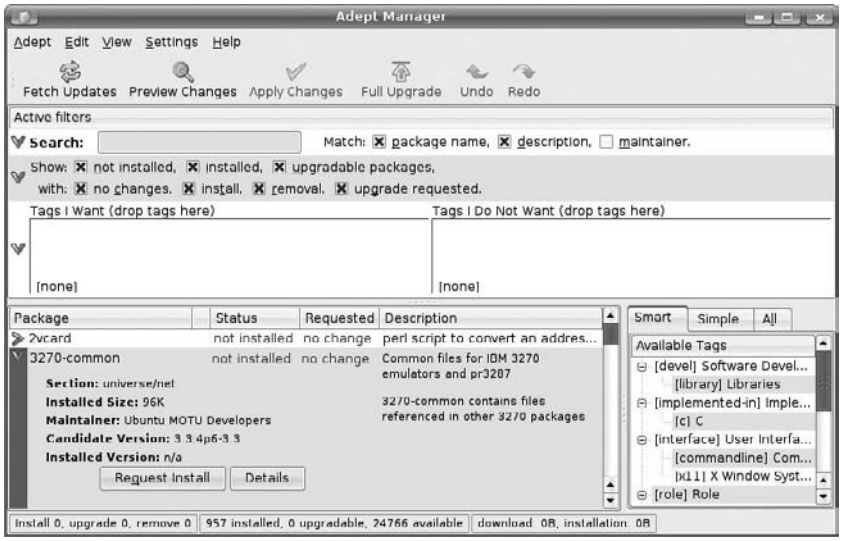


FIGURE 23-33

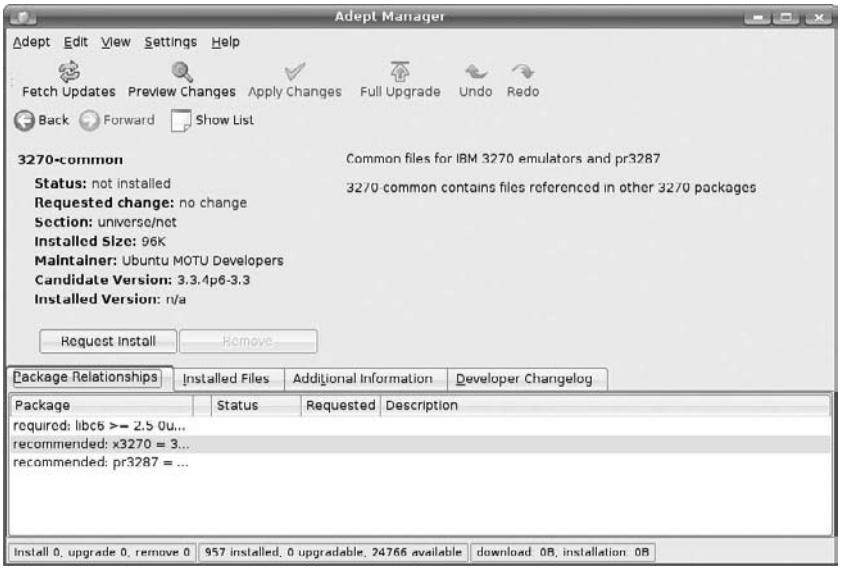
A more verbose package description in the KDE 3 Adept Manager



The expanded description view also provides two buttons: the Request Install button, which marks the current package for installation, and the Details button, which provides even more information about that package and its contents, as shown in Figure 23-34.

FIGURE 23-34

The details view of a package in the KDE 3 Adept Manager



The different tabs on this screen enable you to see the packages that this package depends on (and which will be installed if you select this package for installation and they are not already installed), the files that are installed as part of this package, any additional information that is available about the package, and the log of changes to the package. To close this screen and return to the standard Adept Manager screen, click the Show List button below the toolbar.

- **Tags pane:** Provides three tabs that contain various tags that you can use to refine the list of available packages that are shown in the Package pane. The Smart tab shows tags that are available in the list of packages that is currently displayed in the Packages List pane. The Simple tab shows tags that correspond to those that are available on an Easy tags filter (explained later in this section). The All tab lists all of the tags that are available, regardless of whether any packages that contain these tags are currently displayed.

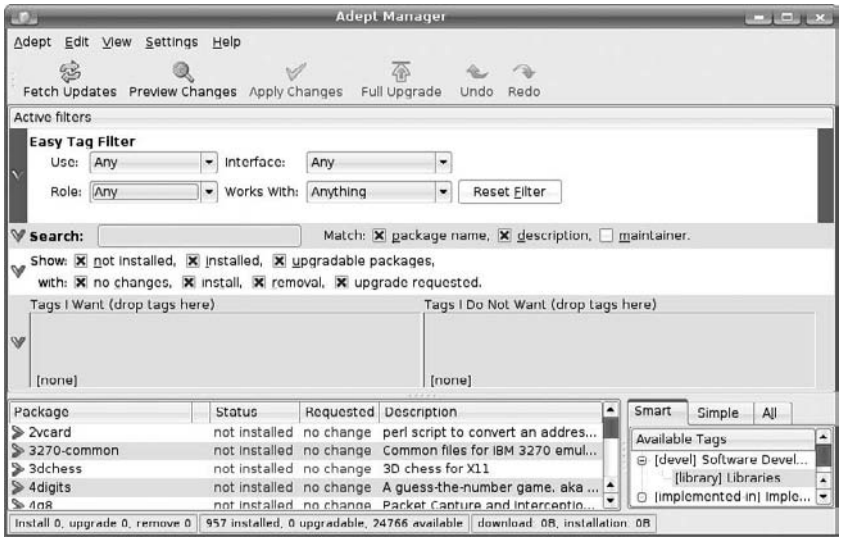
The filters provided in the Active Filters pane make it easy for you to locate a specific package. The following are available filters:

- **Search:** The most commonly used of these is the Search filter, which enables you to search for packages whose name or description (by default) contain the word or words that you enter. You can deselect searching within the package name or description field, or enable searching within the name of the package maintainer, using the checkboxes beside the search filter.

- **Show:** Restricts the list of displayed packages to those with specific values in the Status and Requested fields (as shown in the Packages pane). This is known as a State filter.
- **Tags:** Enables you to restrict package searches and views to packages with specific tags. The Adept Manager takes a slightly different approach to categorizing software than the Synaptic tool does. The bottom-right pane of the Adept Manager dialog provides sets of tags that represent the conceptual categories into which software in your current set of repositories is organized. To restrict your search for packages to a specific category or to exclude packages with certain tags from your search, you can drag and drop tags from the Tags pane into the Tags I Want or Tags I Do Not Want areas.
- **Easy Tags:** Not displayed by default, Easy Tags differ from standard tags in that they provide drop-downs that enable you to limit the packages that are displayed by what they are used for (Use), the type of files that they provide (Role), the type of interface that they use (Interface), and the type of date that they manipulate (Works With). Figure 23-35 shows the Adept Manager with an Easy Tag Filter displayed.

FIGURE 23-35

The KDE 3 Adept Manager with an Easy Tag Filter added



You can add and remove filters by right-clicking in the Active Filters area. This displays a pop-up menu that enables you to remove the filter over which you right-clicked or add another filter of any type.

The Adept Manager doesn't offer many configuration options except for the Settings menu's Configure Shortcuts and Configure Toolbars options, which are quite straightforward and merely define keyboard shortcuts for common actions and change the items for which icons are available on the toolbar. Not very interesting. Let's do some real work…

Searching for Software

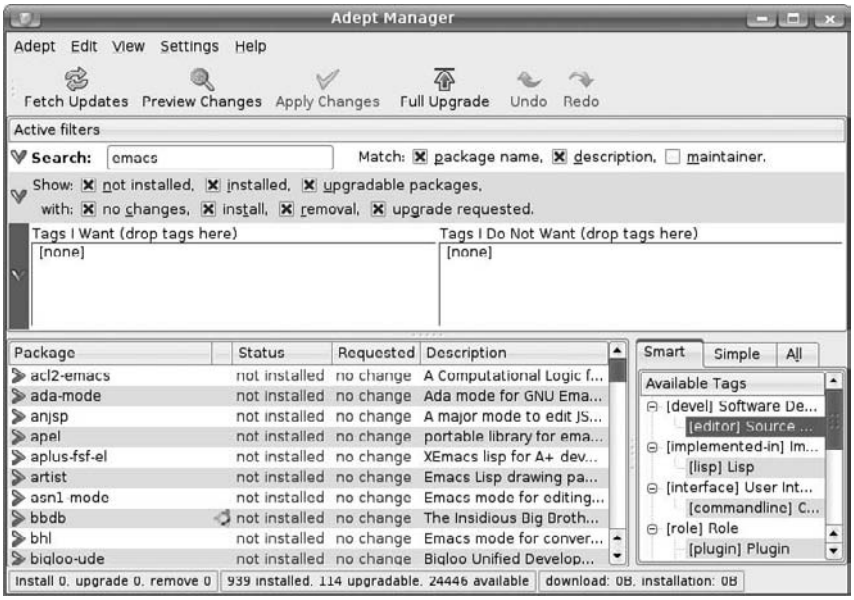
The Adept Manager's Search filter makes it easy to search for packages by some substring of their name without opening a special dialog. One of the hallmarks of the Adept Manager is its immediate usability. To locate packages that are relevant to one or more keywords, enter those terms in the Search box in the Active Filters section of the Adept Manager. Your keyword(s) are matched as you type—typing more specific or additional terms reduces the number of possible matches, but simplifies finding a specific package. By default, the search is done against both the names and descriptions of available packages.

For example, suppose that you're an Emacs text editor fan, have installed the bare minimum number of packages, but are interested in finding out what other Emacs packages are available for use in your development efforts. To find this out:

1. Enter **emacs** as a search term. This restricts the available packages to those that contain the word "emacs" in their name or description. Your Adept Manager screen should look like the one shown in Figure 23-36.

FIGURE 23-36

Listing packages that match "emacs" in the KDE 3 Adept Manager

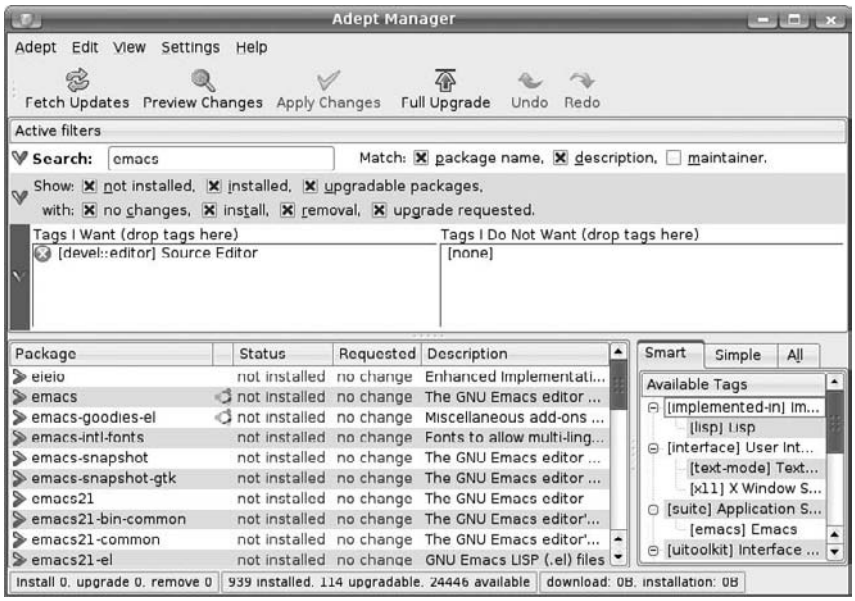


Note that the list of tags in the Tags pane changes as you do a search. This list is automatically assembled from the tags of all of the packages that are currently shown in the Packages List pane.

2. From the Tags pane's Smart tab, select Edit in the [devel] Software Development section, and drag the [editor] tag into the Tags filter's Tags I Want text area. Your Adept Manager screen should look like the one shown in Figure 23-37.

FIGURE 23-37

Refining a package search with a tag in the KDE 3 Adept Manager



Voilà! A much smaller list to search in, which contains only Emacs packages that are related to editing source code.

To clear a filter and start again, simply right-click on any filter that you have changed and select Reset Filter from the context-sensitive menu.

Installing Software

Once you've located a package that you want to install, right-click on its Requested field and select Request Install. The Apply Changes in the toolbar is activated.

If you want to install more than one package, you can use the standard KDE selection conventions of holding down the Ctrl key while clicking the names of individual packages that you want to upgrade—or selecting one, holding down the Shift key, and selecting another, which also selects all of the packages between the two.

To proceed with the installation, click Apply Changes. A progress bar displays as the selected packages are downloaded and installed.

When this process completes, the main Adept Manager dialog displays. You can either continue working in the Adept Manager, or click the Close box or select Adept ➪ Quit to exit the application.

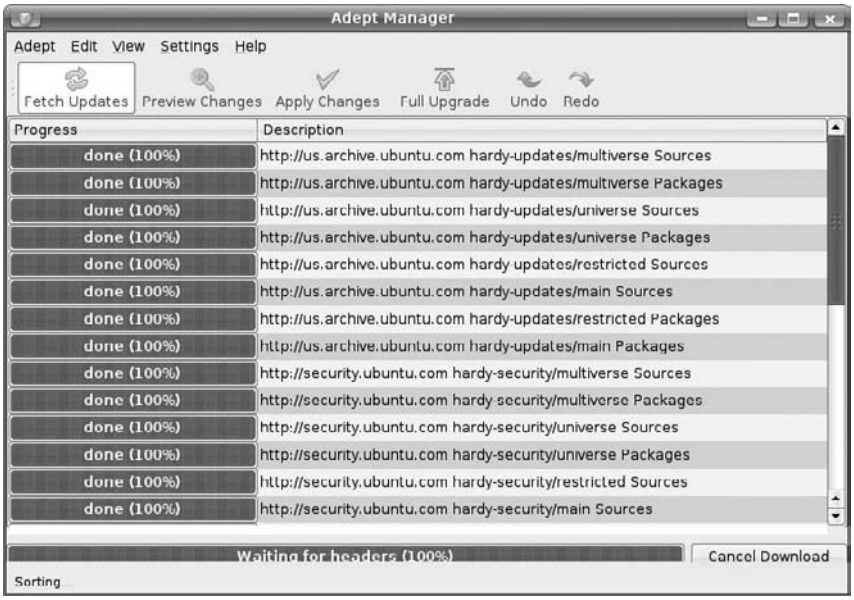
Searching for Software Updates

Although the Adept Updater, discussed later in this chapter, makes it easy to identify and install updated packages, you can also do this from the Adept Manager. The Adept Notifier only searches for new

packages every so often. If you are using the Adept Manager to install new software, you can ensure that you have a list of the latest and greatest packages available by clicking the Fetch Updates button. This displays an Adept Manager screen like the one shown in Figure 23-38 as it retrieves component headers and packages list from the sites listed in your `/etc/apt/sources.list` file.

FIGURE 23-38

Retrieving updated lists of available packages in the KDE 3 Adept Manager



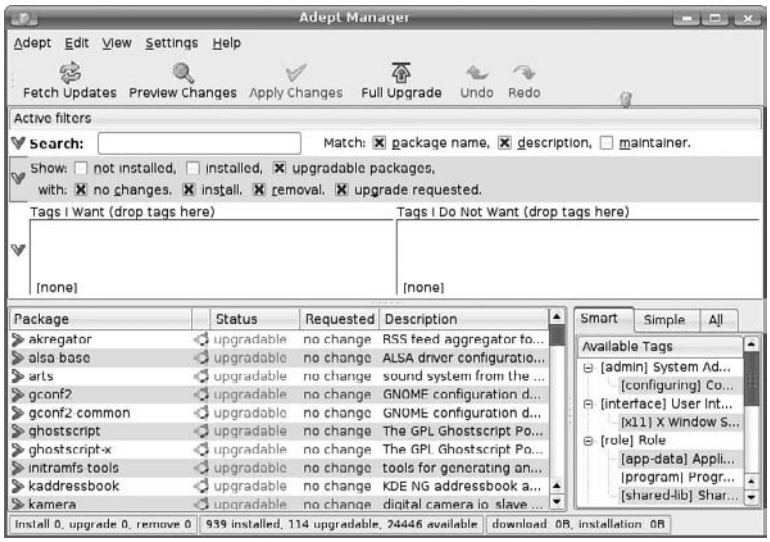
Once this finishes, the middle of the status line at the bottom of the Adept Manager will list the number of installed, updateable, and available packages. If the number of updateable packages is greater than zero, you can see the list by deselecting the not installed and installed status entries in the Show filter. This will display the list of all upgradable packages, as shown in Figure 23-39.

If you're sure that you want to upgrade all of the listed packages, you can simply click the Adept ➪ Safe Upgrade menu item. This menu item is equivalent to the `apt-get upgrade` command or the `aptitude safe-upgrade` commands. If you are sure that no distribution upgrades are available, you could select the Full Upgrade button in the toolbar or the Adept ➪ Full Upgrade menu item, each of which is equivalent to the `apt-get dist-upgrade` or `aptitude full-upgrade` commands. In most cases, you'll want to select the Safe Upgrade command unless you're thinking of upgrading your entire distribution to a newer version that is available.

If you don't want to upgrade all available packages, you can use the standard KDE selection conventions of holding down the Ctrl key while clicking the names of individual packages that you want to upgrade or selecting one, holding down the Shift key, and selecting another, which also selects all of the packages between the two.

FIGURE 23-39

Listing upgradable packages in the KDE 3 Adept Manager



Regardless of how you select them, after selecting the package(s) that you want to upgrade, the Requested field for those packages in the Package List pane will show the word upgrade.

Once you're finished selecting packages, click the Apply Changes button in the toolbar. The updated versions of all of the packages with "upgrade" in their Requested field will be downloaded and installed.

When this process completes, the main Adept Manager dialog displays. You can either continue working in the Adept Manager, or click the Close box or select Adept ➪ Quit to exit the application.

Removing Software in the Adept Manager

Removing software in the Adept Manager is very similar to the installation process. You must first locate the package(s) that you want to remove by searching for their names or scrolling through the Package List pane until you see its name. Either way, once you have located the package that you want to remove, right-click on its name to display the context-sensitive menu.

Marking a package Request Removal removes the package but leaves its configuration files on your system, in case you have already customized them. Marking a package Request Purging removes the package and all of its configuration files.

Once you have selected any packages that you want to remove and marked them for removal, click Apply Changes to proceed with the removal process. A progress screen displays a histogram for each package as it is downloaded, and a final progress bar displays the status of installing all of the downloaded packages.

When this process completes, the main Adept Manager dialog displays. You can either continue working in the Adept Manager, or click the Close box or select Adept ➪ Quit to exit the Adept Updater application.

Using the Adept Updater on KDE 3-Based Systems

Kubuntu's Adept Updater checks the active repositories listed in your `/etc/apt/sources.list` file and notifies you if updates to any installed packages are available. The easiest way to start the Adept Updater is by clicking on the Adept Notifier icon in the system tray, which looks like a small traffic sign with an exclamation mark in the center. Figure 23-40 shows this icon and the tooltip that is associated with the availability of package updates.

FIGURE 23-40

The KDE 3 Adept Notifier icon and tooltip for available updates



Clicking on this icon starts the Update Manager itself, which is shown in Figure 23-41.

Using the Adept Updater is simple—just click Apply Updates. The Adept Updater then automatically retrieves, installs, and configures all available updated packages, displaying progress and status dialogs that are very similar to those used by the Adept Manager. When the update process completes, click Quit to terminate the Adept Updater and begin doing real work.

Using the Adept Manager on KDE 4-Based Systems

The Adept utilities are installed by default as part of any Kubuntu installation. Select the K Menu ➪ Applications ➪ System ➪ Package Manager menu item to start the Adept Package Manager. As when trying to run any graphical administrative application on Kubuntu, you will be prompted for your administrative password before the Adept Manager actually starts. Figure 23-42 shows the dialog displayed when you first start the Adept Manager.

FIGURE 23-41

The KDE 3 Adept Updater

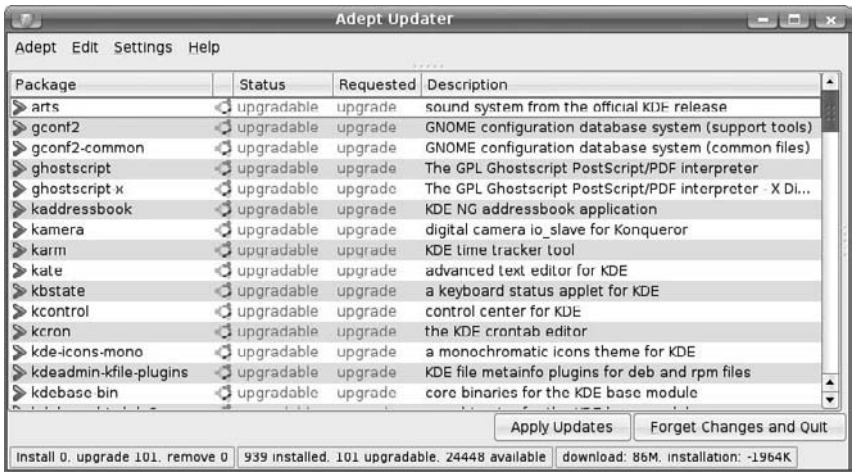
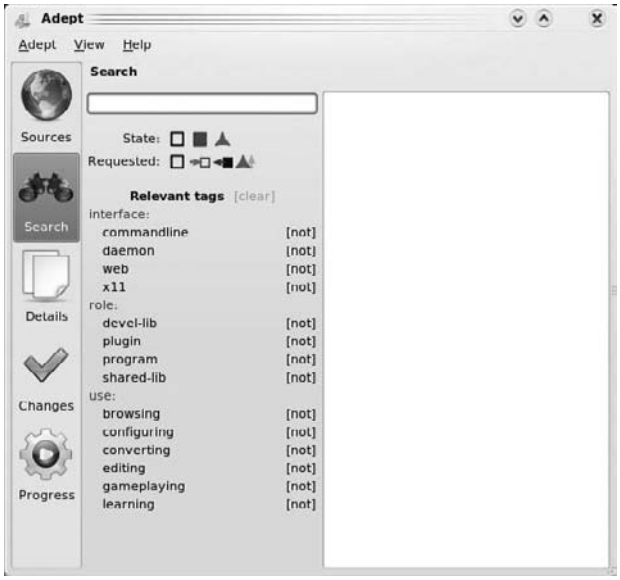


FIGURE 23-42

The startup dialog for the KDE 4 Adept Manager



From top to bottom and left to right, the different areas of the main Adept Manager dialog are as follows:

- **Menus:** Provide easy access to specific Adept commands, views (sections of the Adept manager), and online help.
- **Icon bar:** Provides easy access to the different sections of the Adept Manager, which are the same as the items listed on the Views menu.
- **Search pane:** Displays the filters that you can use to locate a specific package or refine the set of available packages that is currently being displayed. The icons beside the State and Requested labels below the Search text box enable you to filter packages based on different criteria by selecting the associated icon. From left to right, the State filter icons are as follows:
 - **Show available (not installed) packages:** This option actually shows packages that are available on the system, not packages that have not been installed, as the tooltip seems to suggest. The packages listed by this filter are packages that are either already installed or are already installed and are also upgradable.
 - **Show installed packages**
 - **Show upgradable packages**

 From left to right, the Requested filter icons are as follows:
 - **Show packages without any changes requested**
 - **Show packages marked for installation**
 - **Show packages marked for removal**
 - **Show packages marked for upgrade**

 Figure 23-43 shows the Adept Manager after the **S**how available (not installed) packages filter has been selected.

 The Relevant Tags section of the Search pane enables you to refine the list of packages shown in the Package List pane. For example, clicking the x11 filter when you first start the Adept Manager (shown in Figure 23-42) results in the list of packages shown in Figure 23-44, which are all packages that are not installed, but which have been tagged as x11 packages (i.e., graphical packages that require the X Window system, version 11).
- **Package List pane:** Lists the software packages that are available based on the search terms and filters that are currently being used. This area is empty when you first start the Adept Manager, as shown in Figure 23-42. Figure 23-45 shows the Package List pane with a list of packages that match the search term "emacs," which was entered in the Search pane's Search area.

 Clicking on the name of any package in the Package List pane expands the entry for that package, providing detailed information about that package and its contents, as shown in Figure 23-46.

You can see additional information about any package by selecting it in the Package List pane and then clicking the Details icon, which displays a dialog like the one shown in Figure 23-47.

Installing Software

Once you've located a package that you want to install, select its entry in the Package List pane and select the Install This Package checkbox that displays at the bottom of the description of the package. You can then proceed to use search terms and filters to search for other packages that you want to install, selecting the same checkbox to schedule each for installation.

Once you have selected all of the packages that you want to install, click Apply to proceed with the installation. A screen like the one shown in Figure 23-48 displays.

FIGURE 23-43

A list of installed and upgradable packages in the KDE 4 Adept Manager



FIGURE 23-44

A list of packages tagged as x11 that are not installed



FIGURE 23-45

A list of packages that match "emacs" in the KDE 4 Adept Manager

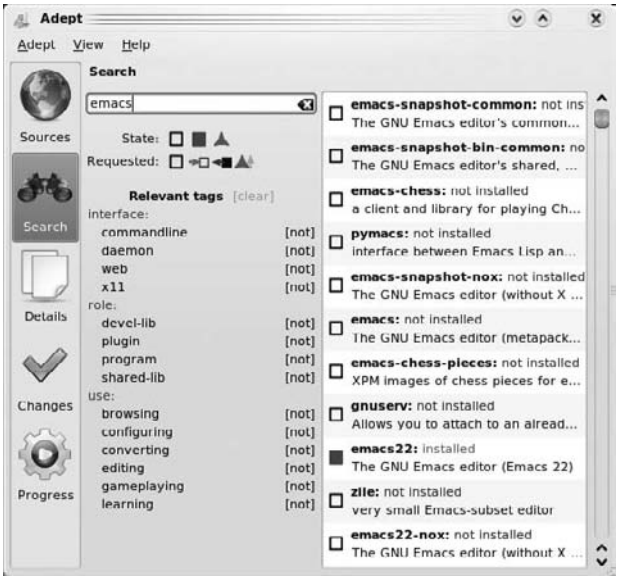


FIGURE 23-46

Expanded information about a package in the KDE 4 Adept Manager

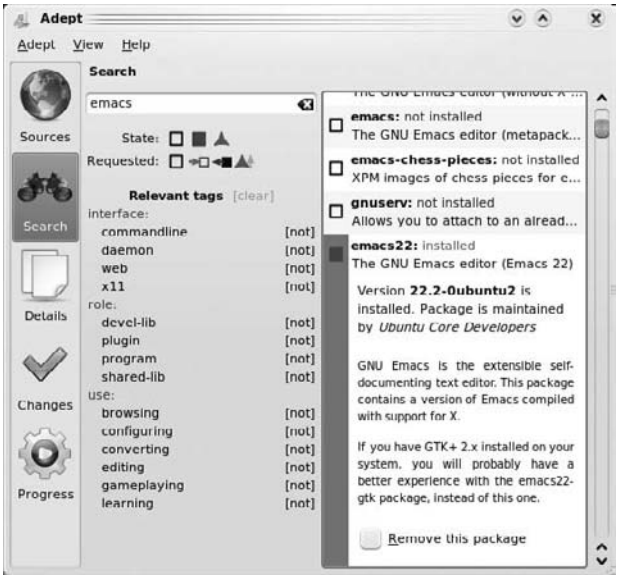


FIGURE 23-47

Detailed information about a package in the KDE 4 Adept Manager

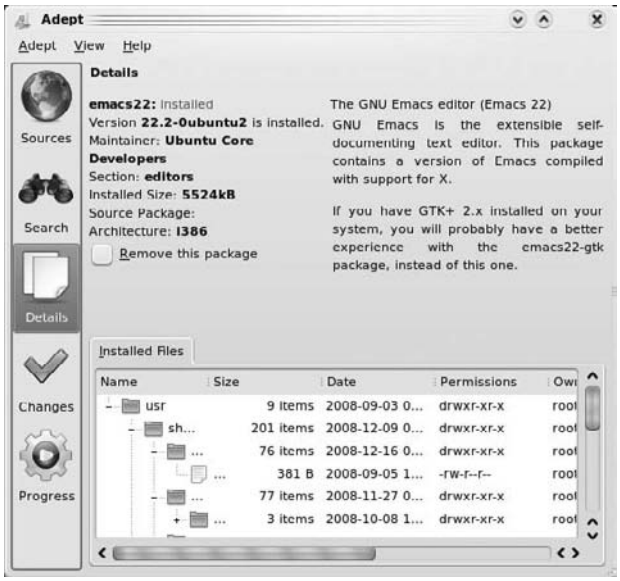
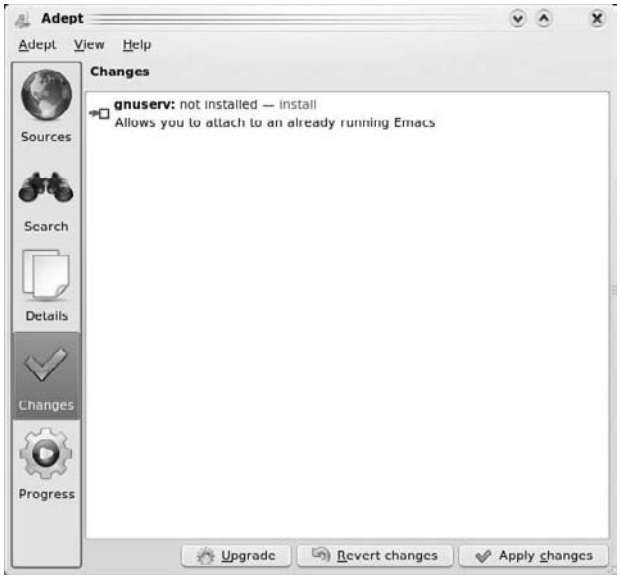


FIGURE 23-48

Packages schedule for installation in the KDE 4 Adept Manager



To proceed with the installation of this package, click Apply Changes. The dialogs displays the installation process. When installation completes, the main Adept Manager dialog displays again. You can either continue working in the Adept Manager, or click the Close box or select Adept ➪ Quit to exit the application.

Searching for Software Updates

Although the Adept Updater, discussed later in this chapter, makes it easy to identify and install updated packages, you can also do this from the Adept Manager. The Adept Notifier applet searches for new packages only every so often. If you are using the Adept Manager to install new software, you can ensure that you have a list of the latest and greatest packages available by clicking the Sources icon and clicking the Fetch current package lists button. This displays an Adept Manager screen like the one shown in Figure 23-49 as it retrieves component headers and packages list from the sites listed in your `/etc/apt/sources.list` file.

FIGURE 23-49

Retrieving updated lists of available packages in the KDE 4 Adept Manager



Once this finishes, you can click the Changes icon and click Upgrade to display the list of all upgradable packages, as shown in Figure 23-50.

If you don't want to upgrade all available packages for some reason, you can select the entry for any package that you do not want to upgrade and clear the Upgrade this package checkbox below its name, which removes it from the list of packages shown in the Adept Manager. Once you're finished with any modifications that you want to make to this list (or if you don't want to make any), click the Apply Changes button. The updated versions of all of the packages that are listed for upgrade will be downloaded and installed.

FIGURE 23-50

Listing upgradable packages in the KDE 4 Adept Manager



When this process completes, you can either continue working in the Adept Manager by selecting the Search icon to look for additional packages that you want to install, or select Adept ➪ Quit to exit the Adept Manager application.

Removing Software in the Adept Manager

Removing software in the Adept Manager is very similar to the installation process. You must first locate the package(s) that you want to remove by searching for their names or scrolling through the Package List pane until you see its name. Either way, once you have located the package that you want to remove, select its name to display more information about the package. You can then select the Remove this package checkbox below its description, click the Apply icon, and click Apply changes to remove that package from your system.

When this process completes, you can either continue working in the Adept Manager by selecting the Search icon to look for additional packages, or select Adept ➪ Quit to exit the Adept Manager application.

Using the Adept Updater on KDE 4-Based Systems

Kubuntu's Adept Updater checks the active repositories listed in your `/etc/apt/sources.list` file and notifies you if updates to any installed packages are available. The easiest way to start the Adept Updater is by clicking on the Adept Notifier icon in the system tray, which looks like a small traffic sign with an exclamation mark in the center. Figure 23-51 shows this icon and the tooltip that is associated with the availability of package updates.

FIGURE 23-51

The KDE 4 Adept Notifier icon and tool tip for available updates



Clicking on this icon starts the Adept Updater, which is shown in Figure 23-52.

FIGURE 23-52

The KDE 3 Adept Updater



Using the Adept Updater is simple—just click Apply changes. The Adept Updater then automatically retrieves, installs, and configures all available updated packages, displaying progress and status dialogs that are very similar to those used by the Adept Manager. When the update process completes, click Adept ➪ Quit to terminate the Adept Updater and begin doing real work.

Converting Packages from Other Package Formats

Although the Ubuntu repositories contain most of the add-on software packages that anyone could ever want, there are cases when you must deal with pre-packed software that is provided in the package formats used by distributions other than Ubuntu. This is often the case with commercial software whose Linux releases are distributed in formats such as RPM and LSB. Luckily, an easy solution to this sort of problem is provided by the `alien` application (`www.kitenet.net/~joey/code/alien.html`),

which is a Perl script that can convert between different package formats. Most important, `alien` can convert most package formats to the DEB package format that is used by the Ubuntu and Debian installation and update tools that are discussed in this chapter. Depending on the format of the package that you are converting from, `alien` may need some additional software to be installed on your Ubuntu system—more about that in a second.

TIP **You must run `alien` as root (preferably using the `sudo` program) so that it can correctly set ownership and file permissions on the files in the converted package.**

The alien application can convert packages to and from the following formats:

- **DEB:** When converting packages to DEB format, the `debhelper`, `dpkg`, `dpkg-dev`, `gcc`, and `make` packages must be installed on your system. (The `gcc`, `make`, and `dpkg-dev` packages are installed as requirements for the build-essential package, as discussed in Bonus Chapter 3. Installing `alien` installs the `debhelper` package.) If you're stuck with using some distribution other than Ubuntu on one or more of your systems, you can use `alien` to convert DEB packages to the appropriate format for your other distribution(s) until you can replace them.
- **LSB:** The package format used to distribute applications and entire distributions that are compliant with the Linux Standard Base specification (`www.freestandards.org/en/LSB`). Because the LSB specification mandates RPM packages, LSB packages are a superset of RPM packages. Therefore, the RPM Package Manager software must be installed on your Ubuntu system in order to convert LSB packages. (The RPM software is automatically installed for you when you install `alien` because it is listed as a requirement for the `alien` package.) If you are converting an LSB package that depends on other packages, the conversion process will correctly translate those dependencies, but will also introduce a new dependency on the `lsb` package, which is available in the Ubuntu repositories. You should also install the `lsb-rpm` package, which is suggested but not required by the `alien` package, if you are converting from LSB packages.
- **PKG:** The package format used by Solaris. Converting PKG files into other package formats requires Solaris-only tools such as `pkginfo` and `pkgtrans`. If you also install the RPM software from one of the Sun Free Software CDs, you can then use `alien` on your Solaris system to convert PKG packages into RPMs. This is probably useful only for packages composed of scripts, source code, or command files because you can't execute Solaris binaries on a Linux system.
- **RPM:** The package format used by distributions such as Red Hat, Fedora Core, Mandrake, Yellow Dog, and so on. To convert packages from the RPM format, the RPM Package Manager, formerly known as the Red Hat Package Manager, software must be installed on your Ubuntu system. (The RPM software is automatically installed for you when you install `alien` because it is listed as a requirement for the `alien` package.)
- **SLP:** The package format used by the extinct Stampede Linux distribution, which ceased development in 2002. I can't believe that this option is still actively supported, but because Stampede Linux is dead, I don't suppose much is changing in its package format.
- **TGZ:** The software distribution format used by most open source software projects and the Slackware Linux distribution, a TGZ file is a `tar` archive that has been compressed using the `gzip` compression tool, whether from the command line or by `tar` itself when creating the archive. Compressed `tar` archives are often simply referred to as *compressed tarballs*. Package files that are generated from TGZ files simply contain the same files as the compressed tarball, plus the package database information for the target package format.

TIP **For more information about Linux package formats and a general comparison of their capabilities, see** `http://kitenet.net/~joey/pkg-comp`.

Using `alien` to convert from one package format to another is easy. The `alien` application provides options such as `--to-deb`, `--to-lsb`, `--to-rpm`, and `--to-tgz` to specify the target output format. The following is an example of converting the RPM package for Adobe Acrobat Reader, which some of you may have heard of, to DEB format:

```
$ sudo alien --scripts --to-deb AdobeReader_enu-7.0.0-2.i386.rpm
Password:
adobereader-enu_7.0.0-3_i386.deb generated
```

Not much to see there, actually. The `--scripts` option tells `alien` to convert any pre- or post-installation scripts found in the package, although the validity of the converted scripts can't be guaranteed—you should watch carefully to make sure that nothing goes amiss when you install the converted package. The `--to-deb` option tells `alien` to generate a Debian package from the input package. If you want to see detailed output, you can add the `-v` option to the `alien` command line, but you're likely to see way more output than you actually care about.

Before installing a converted package, you can verify its integrity using `dpkg`, as in the following example:

```
$ dpkg --info adobereader-enu_7.0.0-3_i386.deb
 new debian package, version 2.0.
 size 37932742 bytes: control archive= 7436 bytes.
     484 bytes,    11 lines      control
   20061 bytes,   207 lines      md5sums
    1340 bytes,    60 lines   *  postinst             #!/bin/sh
     980 bytes,    45 lines   *  postrm               #!/bin/sh
     211 bytes,     7 lines      shlibs
 Package: adobereader-enu
 Version: 7.0.0-3
 Section: alien
 Priority: extra
 Architecture: i386
 Installed-Size: 92316
 Maintainer: root <root@vmdesktop>
 Description: Adobe Reader for Linux. An application that reads a
  PDF document. Adobe Reader 7.0.0 can read documents in PDF format.
  Adobe Reader also allows you to search within PDF files, search
  for PDF files on the internet and participate in collaborative
  document reviews.
  (Converted from a rpm package by alien version 8.64.)
```

This matches quite nicely with information about the original RPM package that was produced on one of my other systems:

```
$ rpm -q --info -v -p AdobeReader_enu-7.0.0-2.i386.rpm
Name        : AdobeReader_enu          Relocations: /usr/local/Adobe/Acrobat7.0
Version     : 7.0.0                         Vendor: Adobe Systems, Incorporated
Release     : 2                     Build Date: Mon 28 Mar 2005 06:34:42 AM EST
Install date: (not installed)            Build Host: acrolinux2.corp.adobe.com
Group       : Applications/Publishing    Source RPM: AdobeReader_enu-7.0.0-2.
    src.rpm
```

```
Size        : 97613446                      License: Commercial
Signature   : (none)
Packager    : Adobe Systems, Incorporated
URL         : www.adobe.com
Summary     : Adobe Reader for Linux. An application that reads a PDF document.
Description :
Adobe Reader 7.0.0 can read documents in PDF format. Adobe Reader
also allows you to search within PDF files, search for PDF files on
the internet and participate in collaborative document reviews.
Distribution: (none)
```

For more information about the `alien` package, see its online reference information, which is available by typing the man `alien` command (after installing the package, of course).

CAUTION **You should never use `alien` in order to install packages containing low-level system utilities or system libraries. Although they may convert and install successfully, these packages will probably not work correctly on your Ubuntu system and you risk reducing your Ubuntu system to a dysfunctional heap of slag that will not boot correctly.**

Keeping Your System Lean, Mean, and Pristine

If you like to play with different software packages as much as I do, you'll find yourself installing lots of random packages on your system, playing with them for a while, and then forgetting about them unless they solve some major problem for you. You may occasionally even be thoughtful enough to remove packages that you aren't planning on using anymore. Unfortunately, this doesn't remove packages that were required, recommended, or suggested by the packages that you've removed (unless you use `aptitude` to both install and remove those package).

Every now and then, it is therefore useful to run a utility that scans all of the packages that are installed on your system and looks for packages that are not used or required by any other package on your system, and which are also not a part of some system installation process. Good examples of programs that you can use to check your system in this way are the `deborphan` and `debfoster` packages. Neither of these packages is installed by default on an Ubuntu system, but you can easily install either of them using your favorite package installation utility.

The core difference between these two packages is that the `deborphan` utility simply produces a list of packages that are not used or referred to by any system package on your system, while the `debfoster` package builds its own database the packages which are required, tries to be smart about things that you may not need, and also asks lots of questions the first time you run it. For all of these reasons, I prefer the `deborphan` package, but I suggest that you investigate both and see which utility best suits your needs and modus operandi.

As mentioned previously, the `deborphan` utility simply lists packages that are no longer required by or associated with any other package. To do something useful with this list, you can redirect it into a file and use that file as the basis for a shell script that will remove the packages that you actually want to remove. Another solution, much more wizardly, is to feed this list to the `apt-get remove` command using the Linux `xargs` command, as in the following example:

```
$ sudo deborphan | xargs apt-get remove --purge -y
```

Before running this sort of command, you should examine the `deborphan` output to make sure that it isn't removing anything that you still want.

Running either the `deborphan` or `debfoster` commands periodically will help you reduce the amount of software package detritus that tends to accumulate on any running system. Keeping your system as free of unneeded packages as possible can help guarantee that the maximum amount of free disk space is available on your system, which you can then devote to "real work" such as your online music archives or collection of artistic digital photographs.

Summary

Package management is a wonderful thing, enabling your system to identify installed packages and, with the help of excellent package formats such as DEB, identify and satisfy any requirements for successfully installing a completely functional software package. This chapter initially discussed the Add/Remove Software menu items that are available on Ubuntu and Kubuntu systems, and explained how to use the associated software (gnome-app-install and Adept Installer, respectively) to locate and install specific software packages. This chapter then introduced the three basic package management utilities that are common to both Ubuntu and Kubuntu systems (`dpkg`, `apt-get`, and `aptitude`), and explained how you can use these and related utilities to install and query packages, figure out other packages that you need, and many more interesting package management tricks.

The primary focus of this chapter was on the graphical installation and automatic software update applications that are available on Ubuntu (Synaptic and Update Manager) and Kubuntu (Adept Manager and Adept Updater) systems. The chapter also focused on how to use the installation utilities to identify and install software packages, and how to use the update utilities to keep your system up-to-date, by installing security fixes or enhanced versions of existing software whenever they are available.

Chapter 24 discusses how to add various types of hardware to your Ubuntu or Kubuntu system. Like most modern personal computer systems, Ubuntu and Kubuntu systems do a great job of automatically recognizing and providing access to new printers, scanners, and removable storage such as USB sticks, USB hard drives, FireWire devices, and so on. Adding new internal drives requires a bit more work, but it is a key element of permanent system upgrades and is easy enough to master.

Chapter 24

Adding Hardware and Attaching Peripherals

Support for a tremendous variety of hardware is a key issue for any desktop computer system. As computer hardware becomes more and more of an end-user commodity, where desktop computer users feel more and more comfortable adding cards and external peripherals to their computer systems, the extent to which an operating system supports all of this new hardware has increasingly higher visibility. An increasing number of devices that are designed as external plug-and-play devices using interfaces like USB and FireWire are available at the local Best Buy, CompUSA, or Stop-n-Go. People expect to buy these things, take them home, plug them in, and find that "they just work." And rightly so.

All hardware requires pieces of kernel software called device drivers in order to be accessed and correctly used by the operating system. This is true for the Microsoft Windows kernel (yes, it has one), the Mac OS X kernel, and the Linux kernel. The fact that CDs labeled "Windows Drivers" fall out of the boxes containing almost every piece of modern hardware might make you think that you'll have a problem using this hardware with Linux. After all, there is no companion coaster labeled "Linux Drivers." What gives?

Luckily, the answer is that the device drivers for most hardware are already available as part of the Linux kernel as installed on your Ubuntu or Kubuntu system. Device drivers are usually either compiled into the kernel or are available as kernel modules that are loaded when specific hardware is detected. (Hence, the generic term for such on-demand drivers: *loadable kernel modules*.) This is true for a vast amount of hardware—after all, in the open source world, all it takes is one motivated Linux user with driver development skills and a piece of hardware that "doesn't quite work right," and a working driver usually surfaces shortly thereafter.

Although I am a fan of Ubuntu and Kubuntu, and a general Linux advocate, I can't deny that the lack of device-specific drivers can sometimes be a problem. To get the absolute best performance out of some pieces of new hardware—most specifically things like graphics and audio cards—you may need special drivers. Unfortunately, most hardware manufacturers are motivated by projected sales numbers based on market analysis, and the Windows market is

IN THIS CHAPTER

- Adding printers
- Adding and using scanners
- Using external drives
- Adding internal drives
- Using PCMCIA cards
- Adding PCI cards

currently at least 50 to a 100 times larger than the Linux market. This is why you get a "Windows Driver" coaster free with almost all new hardware. Open source support for these high-performance drivers may lag behind Windows support, although more and more hardware manufacturers are providing Linux drivers in one form or another, often only available for download from their Web sites. Things are also getting better in that support for the Dynamic Kernel Module Support (DKMS) that was originally introduced by Dell is now present in the Linux kernel, which is a big step in the right direction because it can automatically download and update third-party drivers when you upgrade your kernel.

The good news is that, for the most part, drivers that are required for most new hardware come with your Ubuntu or Kubuntu Linux distribution. Some specialized drivers that are not open source or that have licensing or other encumbrances are available as updates from portions of the Ubuntu repository such as the restricted, universe, or multiverse components, as discussed in Chapter 23. You may have to retrieve other drivers directly from a corporate Web site. The bottom line is that, usually, things will indeed "just work" with Linux.

This chapter explores using different types of hardware devices with your Ubuntu and Kubuntu Linux systems, focusing on the types of hardware that most desktop and laptop computer users will be interested in. Hardware detection and support under Ubuntu, Kubuntu, and Linux in general is constantly improving. Regardless, when researching or buying new hardware, it's always a good idea to spend a bit of time in your favorite search engine, looking for information about Linux support for your hardware. More information is always a good thing, regardless of what the U.S. government says nowadays.

Adding a Printer

Most modern printers connect to your system via its USB interface, although some still connect to the classic parallel port. Regardless of the type of printer you are connecting to your system, the first things to do are to connect it to your system and turn it on. If it is a USB printer, you can verify that the printer is visible to your system by executing the `lsusb` (list USB) command, which provides the information like the following:

```
# lsusb
Bus 007 Device 001: ID 1d6b:0001 Linux Foundation 1.1 root hub
Bus 006 Device 001: ID 1d6b:0001 Linux Foundation 1.1 root hub
Bus 005 Device 003: ID 04b8:0005 Seiko Epson Corp. Stylus D88+
Bus 005 Device 001: ID 1d6b:0001 Linux Foundation 1.1 root hub
Bus 003 Device 002: ID 0d8c:0001 C-Media Electronics, Inc. Audio Device
Bus 003 Device 001: ID 1d6b:0001 Linux Foundation 1.1 root hub
Bus 002 Device 001: ID 1d6b:0001 Linux Foundation 1.1 root hub
Bus 004 Device 001: ID 1d6b:0002 Linux Foundation 2.0 root hub
Bus 001 Device 003: ID 12c8:1f03
Bus 001 Device 001: ID 1d6b:0002 Linux Foundation 2.0 root hub
```

From this output, you can see that an Epson Stylus printer is attached to this system via USB.

TIP **If you have not purchased a printer yet and want to get information about Linux support for a specific printer, check out the LinuxPrinting.org site at `www.linuxprinting.org`. This site, which actually resolves to a site at the Linux Foundation (`www.linuxfoundation.org/en/OpenPrinting`) provides an easy interface at `www.openprinting.org/printer_list.cgi` that enables you to get information about how well specific printers (or all of the printers from a specific manufacturer) work with the Linux printing system. For more information about printing on Linux, see Bonus Chapter 12, "Setting Up a Print Server."**

Adding a locally attached printer has become incredibly simple thanks to the fact that most printers use USB interfaces nowadays rather than the classic Centronics parallel interface that many of us remember with varying degrees of fondness. Because attaching a USB device to your system causes a recognizable event that your system can detect, your system can also initiate its configuration. The next few sections discuss what happens when you attach a new printer to your Ubuntu or Kubuntu system, and how you can create or customize that printer's configuration, if necessary.

Adding a Printer on Ubuntu

Earlier releases of Ubuntu enabled you to configure what should happen when a new printer was detected by your system. This was friendly but largely a waste of time because the answer is almost always "Please configure the thing so that I can actually print on it." On a Ubuntu 8.10 system, attaching a USB printer usually results in a pop-up like the one shown in Figure 24-1—the printer is recognized, added, and configured, and your system is ready to use it.

FIGURE 24-1

Auto-detection and configuration of a USB printer



If the printer was recognized correctly, your new printer is automatically added to the list of printers that are available on your system, and you can simply click the red x to close this dialog. You're ready to print! However, if your printer was not recognized correctly or if you simply want to test or further configure it, you can click Configure to open the `system-config-printer` application, which is discussed in the next section.

Unfortunately, only USB printers and shared networked printers can be recognized automatically by your Ubuntu system. If you're adding a printer that uses the parallel port or need to define a networked printer that your system can't detect automatically, skip ahead to the section entitled "Manually Defining a New Printer."

Configuring an Existing Printer

You can fine-tune the settings for any printer on your system at two points: by clicking the Configure button on the dialog that appears when you first attach your printer, as introduced in the previous section, or by selecting the System ➪ Administration ➪ Printers menu item. Either of these options displays the general `system-config-printer` dialog shown in Figure 24-2.

If you click Configure when you are connecting a printer for the first time, the configuration dialog for your new printer also appears; it looks look like the dialog shown in Figure 24-3. If you are configuring a printer after the initial installation process has completed, you can display this dialog by right-clicking on the printer that you want to configure and selecting Properties from the context-sensitive menu that is displayed.

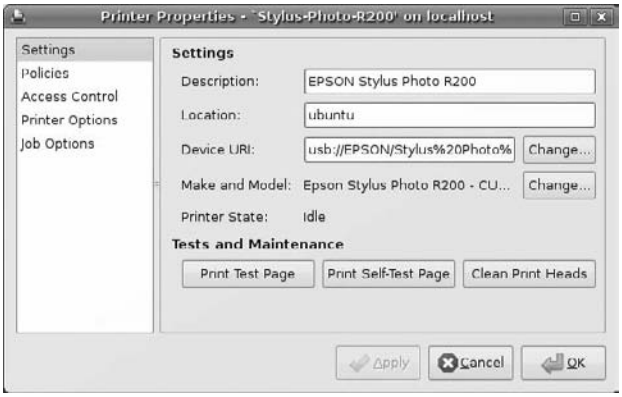
FIGURE 24-2

The Ubuntu Printers utility



FIGURE 24-3

The configuration dialog for a specific printer



This dialog lists a number of different configuration options in its left pane. The Settings pane, which is selected by default, enables you to print a test page to verify that your printer is working correctly, identify how the printer is connected to your system, change the printer type if it was recognized incorrectly, and perform any printer-specific tasks that are available, such as cleaning the print heads on inkjet-style printers.

If your printer was recognized incorrectly or if you want use drivers that have come with the printer, click Change to display the Change Driver dialog. With the exception of the window title, this dialog is identical to one of the dialogs used when configuring a new printer (shown in Figure 24-5), which is discussed in the next section.

The other configuration options available in this dialog do the following:

- **Policies:** Enable you to control the state of the printer (is it enabled and accepting jobs, and is it shared so that other systems can access it?), general policies such as what happens if a printing error occurs, and banner options that enable you to print cover pages that can precede and follow each print job with various banners, including banners that identify print jobs as Classified or Top Secret so that any spies in your office will know which documents to steal.

- **Access Control:** Enables you to configure who can print on that printer, either by disallowing printing by everyone but selected users, or by allowing printing by everyone except for selected users.
- **Printer Options:** Lets you specify the paper size, type of paper being used, the paper tray from which to obtain paper (manual feed, paper tray, and so on), the level of print quality to use and the resolution in which to print, and so on. This dialog also enables you to specify whether you want the printer to print in grayscale or color and, if in color, which color model to use (RGB, CMYK, and so on). I know what you're thinking, but unfortunately this option cannot be used to enable you to print in color on a black and white or grayscale printer.
- **Job Options:** Lets you set defaults for the number of copies of each document to print, the default font size and margins for generic text documents, the page orientation, number of pages per side, scaling, and so on.

After changing any of these options, click OK to save your changes and close this dialog. The main printer configuration utility dialog, shown in Figure 24-2, will still be visible on your screen.

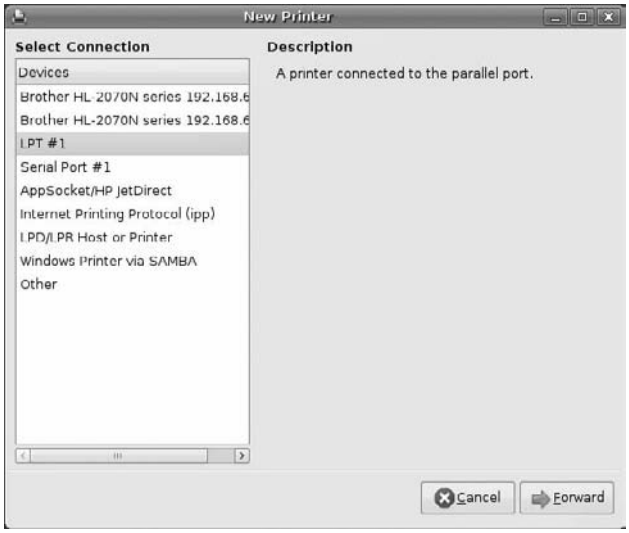
Manually Defining a New Printer

The dialog shown in Figure 24-2 enables you to configure existing printers or manually add new printers to your system, including shared printers. You can start this utility after attaching a new printer (as explained previously) or at any time by selecting the System ➪ Administration ➪ Printing menu command.

To begin manually defining a new printer, click the New Printer icon or select Server ➪ New ➪ Printer to start the wizard that helps configure a new printer. After the Printers utility processes and parses its database of available printers, the first page of the New Printer dialog is displayed, as shown in Figure 24-4.

FIGURE 24-4

Specifying how a new printer is connected.



This dialog displays various choices for how your system will communicate with the new printer. Depending on the connection mechanism that you select, you should be able to obtain all of the necessary information from the documentation for your printer, from your local system administrator or by examining your system. Here are the connection options:

- **LPT #1:** A printer that is connected to your system's parallel port.
- **Serial Port #1:** A printer that is connected to a serial port on your system. Serial printers are extremely uncommon nowadays, but if you are using one, you will need to provide information about your new printer. Enter the communication speed (Baud rate) that your new printer supports, and change the type of parity, the number of data bits, and the way in which the printer tells the system to start/stop sending data (Flow Control). The default values (8 data bits, no parity, and one stop bit) are usually the correct settings.
- **AppSocket/HP JetDirect:** A printer that you can print directly with over the network by using a specific IP address and port. This is a common way of printing to modern standalone, network-aware printers from manufacturers such as Brother, HP, Epson, and so on. After selecting this option, you will need to provide information about your new printer. Enter the name or IP address of the printer in the Host field. You can leave the default port number (9100) unchanged unless you know that your printer uses a different port for JetDirect connections.
- **Internet Printing Protocol (ipp):** A printer whose queue is managed by a specific host. IPP is a common alternative to JetDirect but supports capabilities such as authentication, access control, and host-side document formatting, and printer management. You will need to know the name or IP address of the host that manages the printer and the name of the queue for the specified printer on that host. You can use the Find Queue button to list all of the publicly available queues on a specified host.
- **LPD/LPR Host or Printer:** A printer whose queue is managed by a specific host (including the local host) using the traditional UNIX Line Printer Damon (LPD) protocol. You will need to know the name or IP address of the host that manages the printer, and the name of the queue for the specified printer on that host. You can use the Probe button to list all of the publicly available queues on a specified host.
- **Windows Printer via SAMBA:** A printer that is managed by a host located on a Microsoft Windows network. Such hosts are typically Windows servers, but can also include Samba servers running on other UNIX or Linux systems on your network. You will need to know the name of your local workgroup or domain, the name or IP address of the host that manages the printer, and the name of the queue for the specified printer on that host. You can use the Browse button to navigate through your local workgroup/domain and hosts to locate the publicly available queues on a specified host.
- **Other:** A printer that you can access through a Universal Resource Identifier (URI) that is used to uniquely identify a resource on the Internet.

Select the mechanism that your system uses to connect to the printer that you are adding. Click Forward to proceed to the dialog shown in Figure 24-5.

The dialog shown in Figure 24-5 enables you to provide your system with the information that it requires to format documents in the manner required by your new printer. This dialog provides three ways to identify your printer:

- **Select printer from database:** The default option, this mechanism enables you to identify the driver(s) for your printer by identifying the vendor of your printer and then identifying its model

number in a subsequent dialog. The combination of your printer vendor and model number enables the printer configuration tool to select (or create) an appropriate PostScript Printer Definition (PPD) file for your printer.

- **Provide PDD file:** PPD files are typically provided on a CD that accompanies your printer, and provide drivers and related information that describe how to print to your new printer. If you received a CD or other media with your printer, you should try this option first because the PPD files that are delivered with a printer typically provide the most up-to-date information about how to print to a specific printer. If you have trouble finding the PDD file for your printer, it may be located in the directory where Microsoft Windows printer driver(s) are located.
- **Search for a print driver to download:** Selecting this option causes the printer configuration utility to search the `www.openprinting.org` site for an appropriate driver for your printer. If multiple possibilities are located, you can select among them from the drop-down Printer model list.

FIGURE 24-5

Identifying your printer

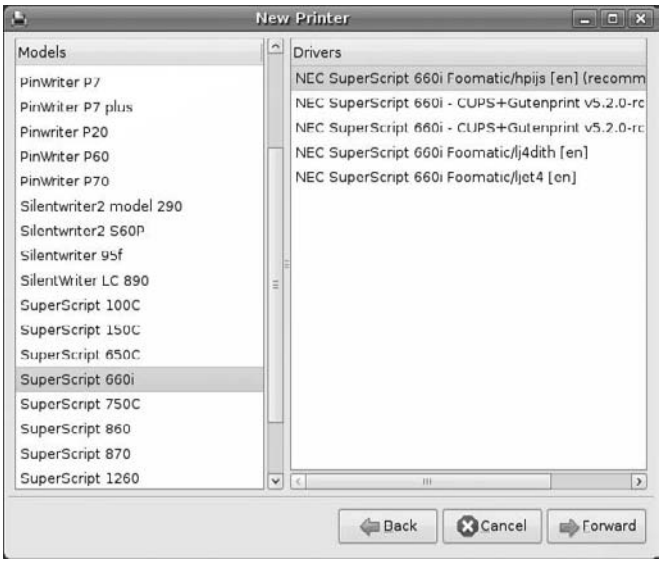


After making your selection, click Forward. The dialog that you see next depends on your selection on the current dialog:

- **Select printer from database:** Clicking Forward after specifying this option and selecting the manufacturer of your printer displays the dialog shown in Figure 24-6.
- **Provide PDD file:** Clicking Forward after specifying this option and selecting the PPD file for your printer takes you to the dialog shown in Figure 24-8.
- **Search for a print driver to download:** Clicking Forward after specifying this option and selecting the PPD file for your printer takes you to the dialog shown in Figure 24-7, which will list any downloadable PPD files that are available.

FIGURE 24-6

Selecting the model of your printer.



The dialog shown in Figure 24-6 appears only if you are manually specifying the manufacturer and model of your printer to locate the correct printer driver. Locate the model of your printer or, if it is not specifically listed, locate the model number that is closest to your printer's. For example, I own a Brother HL-2070N printer, which was not listed at the time that I updated this book, but the driver for the Brother HL-2060 printer works fine with mine. After locating and selecting the best printer driver for your printer, you may be able to select different drivers for that printer in the driver list at the right of this dialog. The driver identified as recommended is usually the best one to use—you can always edit your printer (as described in the previous section) if the driver that you initially select doesn't work correctly for you.

After selecting the best printer driver for your printer, click Forward to proceed. The dialog shown in Figure 24-8 appears.

The dialog shown in Figure 24-7 appears only if you have manually specified the location of a PPD file that specifies the driver for your printer. In many cases, such as the example shown in the figure, no additional download is required—an appropriate PDD file for your printer is included in your Ubuntu installation. In this case, click Back to return to the previous dialog, select the Select printer from Database option, and locate the driver for your printer in that fashion. If a different PPD file for your printer is available for downloading from the `www.openprinting.org site`, select it from the list at left and click Forward to download that driver. The dialog shown in Figure 24-8 appears after the download completes.

The dialog shown in Figure 24-8 enables you to provide some meaningful (but optional) information about your printer and to optionally change the name assigned to that printer. Once you are happy with the values on this screen, click Apply to complete the definition for your new printer. This closes the printer configuration dialog and displays an icon for your new printer in the dialog shown earlier in Figure 24-2.

FIGURE 24-7

Downloading or installing a PPD file



FIGURE 24-8

Naming your printer



To test your new printer, right-click on its name and select Properties from the pop-up menu. This displays a Properties dialog like the one you saw in Figure 24-3. To verify that your printer is working correctly, click Print a Test Page. This queues and prints a test page for your printer. If the test page does not print, use the other tabs in the Properties dialog to verify that the correct connection settings are being used, and that the right printer manufacturer and driver for your printer are specified. You should then be able to successfully print a test page.

TIP **If your printer is not automatically detected or cannot be made to print to with any manual setting specified in the Add printer dialog, try rebooting your system with the printer connected and turned on. In some rare cases, Ubuntu's printer and device subsystems may not correctly detect a device unless it is present and active when you boot your system.**

Adding and Configuring Printers on Kubuntu

The transition from KDE 3 to KDE 4 has led to some changes in how you add and configure printers, as well as in the degree to which they can be configured. The processes are largely the same across the two KDE versions, although the KDE 4 dialogs have simplified the process. On the downside for KDE 4 fans, the KDE 3-based Kubuntu 8.04 and earlier systems have some configuration bells and whistles that haven't made it the KDE 4-based Kubuntu 8.10 and later systems at the time that this book was updated. For the most part, these are truly bells and whistles—and the list of KDE 4 features is growing longer every day.

Adding Printers on KDE 3-Based Kubuntu Systems

If you are using a Kubuntu 8.04 or earlier system, adding a new printer or reconfiguring an existing one is done through the K Menu ➪ System Settings application. After starting this application, selecting the Printers icon, located in the Computer Administration section, displays a screen like the one shown in Figure 24-9.

This dialog enables you to configure and manage printers and related output devices such as printing to PDF or PostScript output files, or printing to Fax devices.

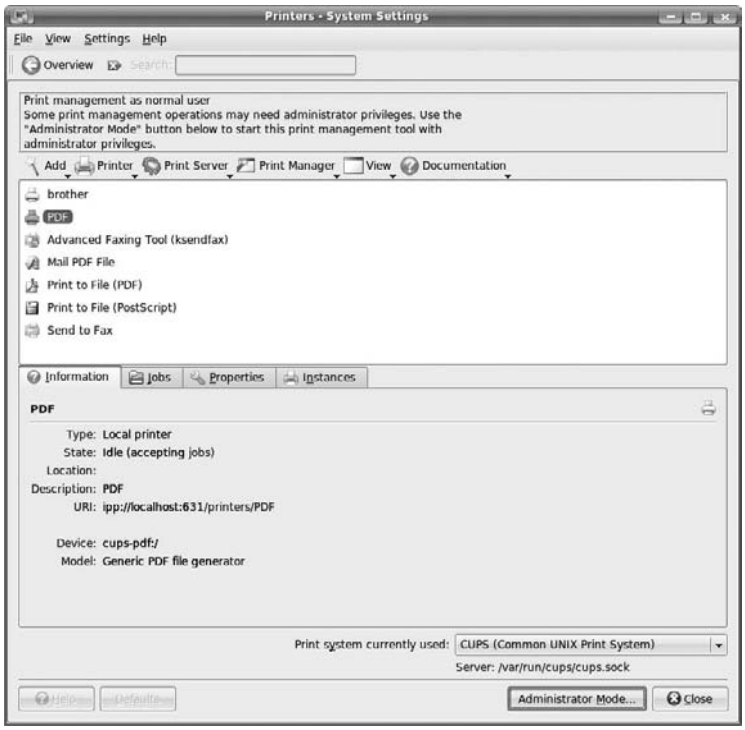
To add a new printer, connect the printer to your system and turn it on. KDE 3-based Kubuntu systems automatically detect and configure most modern USB printers—if you're lucky, you will see two pop-up text balloons in the lower-right corner of your screen after turning on your USB printer. The first will note that the printer has been detected and that support for that printer is being configured. The second, shown in Figure 24-10, tells you that configuration has been successful. In this case, you can immediately print to your new printer, and you can relate the ease of the whole "new printer" experience on Kubuntu to your Windows and Mac OS X friends whenever the subject comes up in some future nerd conversation.

If you are using a new printer that cannot be automatically detected as new hardware, KDE provides a friendly Printer Wizard that walks you through the steps for identifying and configuring the printer. Printers that you will have to manually configure include those that are connected to your system's parallel port and most types of network printers.

To start the Printer Wizard, click Add in the Printers dialog and select Add Printer/Class from the drop-down menu that appears. The Printer Wizard starts, displaying a welcome dialog. Click Next to proceed to the actual printer configuration process, which beings with the dialog shown in Figure 24-11.

FIGURE 24-9

The KDE 3 Printer configuration applet's initial screen



From the following, select the option that identifies how your system will connect to the new printer:

- **Local printer (parallel, serial, USB):** A printer that is directly connected to your system through a local interface such as a parallel or serial port, or a USB connection.
- **Remote LPD queue:** A printer whose queue is managed by a specific host (which can be the local host) using the traditional UNIX Line Printer Damon (LPD) protocol.
- **SMB shared printer (Windows):** A printer that is managed by a host located on a Microsoft Windows network. Such hosts are typically Windows servers but can also include Samba servers running on other UNIX or Linux systems on your network.
- **Network printer (TCP):** A printer that you can print directly to over the network by using a specific IP address and port. This is a common way of printing to modern standalone, network-aware printers from manufacturers such as Brother, HP, Epson, and so on.
- **Remote CUPS server (IPP/HTTP):** A printer whose queue is managed by a specific host. The Internet printing protocol (IPP) is a common alternative to network printing mechanisms such as JetDirect but supports capabilities such as authentication, access control, and host-side document formatting and printer management. CUPS (Common UNIX Printing System), which is discussed in Bonus Chapter 12, is a specific implementation of an IPP print server that implements printing services using HTTP as a transport protocol.

- **Network printer w/IPP (IPP/HTTP):** A printer whose queue is managed by a specific host. This option covers non-CUPS print servers that still implement IPP using HTTP as a transport protocol.
- **Other printer type:** A networked printer that you can identify by specifying its URI (Uniform Resource Identifier).

FIGURE 24-10

Automatic printer configuration on a KDE 3-based system

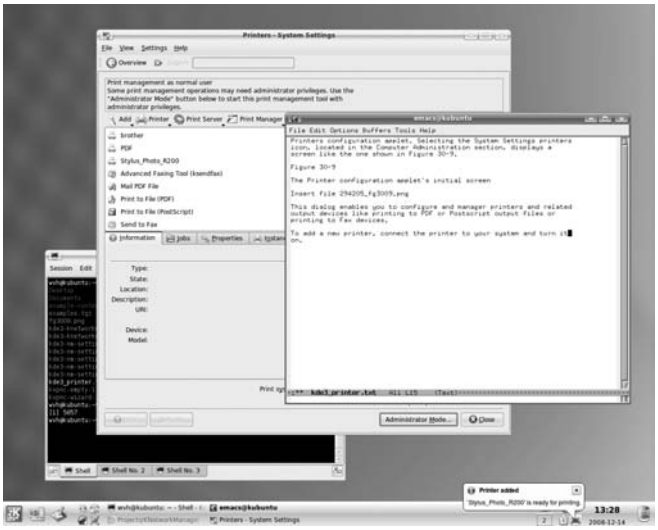
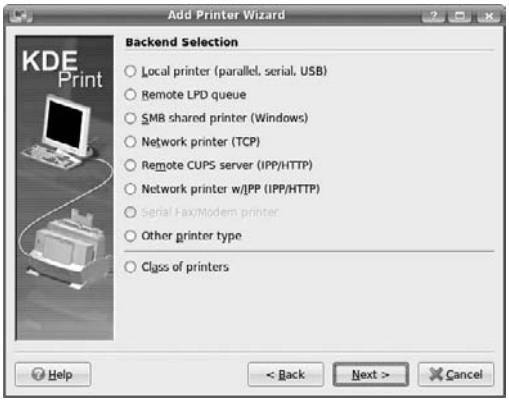


FIGURE 24-11

Identifying how you connect to your printer



Click Next to proceed. The dialog that appears next depends on the connection option that you selected in Figure 24-11:

- **Local printer (parallel, serial, USB):** A dialog enables you to select how your printer is physically connected to your Kubuntu system. For older printers that use a classic Centronic printer interface, this will be your system's parallel port. USB printers that your system can't recognize as a specific model will be listed in the USB or Others sections.
- **Remote LPD Queue:** The dialog provides text fields in which you can enter the name or IP address of the system that provides access to the printer and the name of the queue for that printer on the specified system.
- **SMB shared printer (Windows):** A dialog enables you to specify the account that you will use to connect to the Microsoft Windows system to which the printer is actually attached. (This can also be a Linux system that exports its printers via Samba.) Your choices are using an Anonymous, Guest, or Normal account. If you select the Normal account radio button, you will need to enter the login and password for a valid account on the remote system.
- Clicking Next after entering authentication information displays a dialog that enables you to specify the workgroup/domain of the system that is hosting the printer that you want to connect to, the hostname or IP address of the server to which the printer is connected, and the name of the printer on that system. If you don't know all of this information offhand, you can click Scan to probe your local network and locate available printers.
- **Network Printer (TCP):** The dialog provides text fields in which you can enter the IP address (or hostname) of the printer that you want to connect to and the port on which the printer is listening for network printing requests. This is typically port 9100 for networked printers that use the HP DirectJet printing mechanism. If you don't know this information offhand, you can click Settings to enter the base IP address of your local network, and then click Scan to probe your local network for available networked printers. Note that scanning may not find all of your networked printers if they automatically put themselves to sleep after a given period of time.
- **Remote CUPS server (IPP/HTTP):** The dialog enables you to specify the account that you may need to use to connect to the system to which the printer is actually attached. Your choices are using an Anonymous or Normal account. If you select the Normal account radio button, you will need to enter the login and password for a valid account on the remote system.
- Clicking Next after entering authentication information displays a dialog that enables you to specify the hostname or IP address of the server to which the printer is connected, and the port on which the print server on that system is listening.
- Clicking Next after entering network address and port information displays a dialog that lists all of the printers that the print server at the specified address/port knows about.
- **Network printer w/IPP (IPP/HTTP):** The dialog enables you to enter the URI (Uniform Resource Identifier) that you will use to connect to the networked printer. If you don't know this information offhand, you can click Settings to enter the base IP address of your local network, click Close to close the Settings dialog, and then click Scan to probe your local network for available IPP print servers. After selecting an IPP print server, you may also have to append the name of the printer to which you want to print to the default URI that this dialog provides.
- **Other printer type:** The dialog provides a text area in which you can enter the URI for a networked printer resource. The bottom portion of this dialog displays a list of all networked and local printers that could be detected.

After providing the information that is required by the connection that your printer uses, click Next to proceed. The dialog shown in Figure 24-12 appears.

In this dialog, you select the manufacturer of your printer from the list at left and the specific model of your printer from the list at right. If you don't see an entry for the manufacturer of your printer and your printer advertises compatibility with printer languages such as HP PCL or PostScript, select GENERIC as the manufacturer of your printer. If your printer manufacturer is listed, but the exact model of your printer isn't present in the list, try selecting another model of printer from the same vendor that is as close as possible to yours. If you don't see a remotely similar model and your printer advertises compatibility with another type of printer, you can select the manufacturer and model of that printer. For example, printers that advertise HP LaserJet compatibility are almost always compatible with a LaserJet 3 or LaserJet 4 model. Similarly, printers that advertise PostScript compatibility can usually be configured by selecting the PostScript printer checkbox on this dialog, which selects a generic PostScript driver.

After selecting the appropriate Manufacturer and model for your printer, click Next to proceed. A dialog like the one shown in Figure 24-13 appears.

In this dialog, you select a specific print driver to use with your new printer. The recommended option is preselected, and is usually the right one to try. You can always modify the printer later and use a different driver if the suggested one doesn't work optimally for you. After making your selection, click Next to proceed. The dialog shown in Figure 24-14 appears.

A banner is a separate page that precedes and/or follows each of the documents that you print to the specified printer. These are often referred to as job or page separators because they separate each print job from the next, and make it easy to identify all of the pages associated with a specific document and determine if your entire document has finished printing. One the other hand, they are a significant waste of paper. A starting banner is usually sufficient in enterprise or academic environments where many users are printing to the same printer.

FIGURE 24-12

Selecting your printer's manufacturer and type

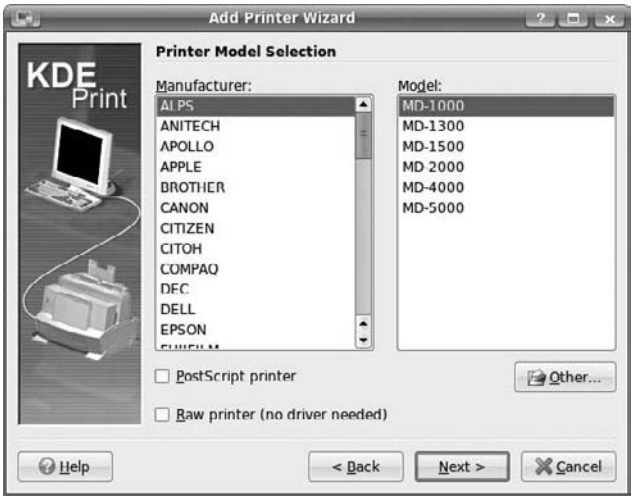


FIGURE 24-13

Selecting a driver for your printer



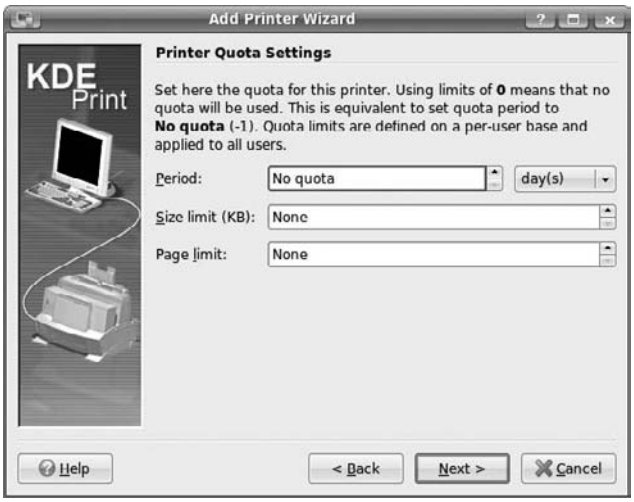
FIGURE 24-14

Adding a starting or ending page banner



FIGURE 24-15

Defining print quotas



After making any changes, click Next to proceed. The dialog shown in Figure 24-15 appears.

The dialog enables you to limit the use of your new printer, restricting the size or number of pages that users can print on it over a specified number of minutes, hours, days, weeks, or months. The default selection, No quota, means that anyone can print whatever they want on your printer at any time. Unless carefully specified, imposing print quotas can prevent you or other users of your printer from printing things that you actually need to print. However, they can be useful to restrict access to specific printers (such as color printers) or to prevent coworkers from printing 200 copies of their personal holiday newsletter in enterprise or academic environments.

After specifying any restrictions that you want to impose, click Next. The dialog shown in Figure 24-16 appears.

This dialog enables you to limit access to your new printer. If you want to prevent people from printing to your new printer, you can either do so by enabling for everyone and denying access to specific users (which is the default) or by denying access for everyone and enabling access for specific users. This option can be quite useful for preventing people from abusing expensive color or extremely high-quality printers.

To disable or enable access by specific users, select Denied Users or Allowed Users, and enter the login names of any users for whom you want to deny or enable access.

Once you're done making changes to this dialog (or if you don't want to change anything), click Next. The dialog shown in Figure 24-17 appears.

You must supply a name for your new printer because the name is used for system resources, such as print queues, that your system will use to manage access to the printer. A location string is optional but can be handy to help other users of your printer to identify a specific printer (assuming that you have more than one).

FIGURE 24-16

Restricting access to a printer



FIGURE 24-17

Naming your new printer



After specifying the name of your printer, providing an optional location, and optionally modifying its description, click Next. A summary dialog displays all of the information that you provided when

configuring your new printer. After reviewing this information, click Finish to accept the configuration for your new printer. To change anything, click Back multiple times to return to the dialog that contains the settings that you want to change, make those changes, and then click Next repeatedly to return to this dialog and accept those changes.

At this point, you should be able to print to your new printer from any application on your system. To test your new printer, right-click the name of your new printer in the Printers dialog, and select Test printer from the context-sensitive menu that displays. A confirmation dialog displays, verifying that you are about to send a test page to the selected printer. Click Print Test Page to do so.

After you are finished defining and testing your printer, you can click Back to return to the System Settings application, or click the close box or select File ➪ Quit to exit the Systems Settings application all together.

Configuring an Existing Printer on a KDE 3-Based System

In most cases, the settings for the printers that you have configured on your system are the ones that you actually want to use. However, over time, you may find that you want to change some aspects of how you and other users of your systems interact with your printer. For example, after losing too many print jobs to other people who have accidentally taken them away, you may want to add a print banner before each print job. Similarly, if you're always running out of ink because an amateur photographer in your department (or even at home) uses your printer instead of the local Foto-Hut or Copy-Quick, you may want to restrict access to your color printer so that you can actually use it occasionally.

In addition to general administrative measures like these, you may also need to fine-tune certain aspects of your printer such as the print quality, type of paper, the order in which your documents print, and so on. KDE provides different mechanisms for changing general administrative options for your printer or to fine-tune its configuration. Both of these options are provided by the Printers applet, which you can start by selecting K Menu ➪ System Settings ➪ Printers icon. Once the Printers window appears, select the printer that you want to modify. Next:

- To change administrative options that you specified when initially configuring your printer, click the Properties tab in the lower portion of the Printers dialog to display information about the selected printer, as shown in Figure 24-18.
- Click the icons in the list at the left of the bottom half of this dialog to review the current settings for your printer, and click Change to pop-up a dialog that enables you to reconfigure that aspect of the general administration of your printer. (These are almost identical to the dialogs that you saw when initially configuring that aspect of your printer.) After making your changes in that dialog, click Finish to close the dialog and save your changes, or click cancel to close the dialog without saving your modifications. The Printers dialog is redisplayed.
- To fine-tune your printer, right-click on the name of that printer in the top portion of the dialog shown in Figure 24-18, and select Configure from the context-sensitive menu that displays. A dialog like the one shown in Figure 24-19 appears.
- To modify any of the settings on this dialog, select that setting in the top portion, which displays the available settings for that option in the bottom portion. For example, Figure 24-19 shows the possible Media Source settings (that is, what you're printing on) for one of my printers.
- After making your changes, click OK to close the dialog and save your changes, or click Cancel to close the dialog without saving your modifications. The Printers dialog appears.

FIGURE 24-18

Modifying administrative printer properties

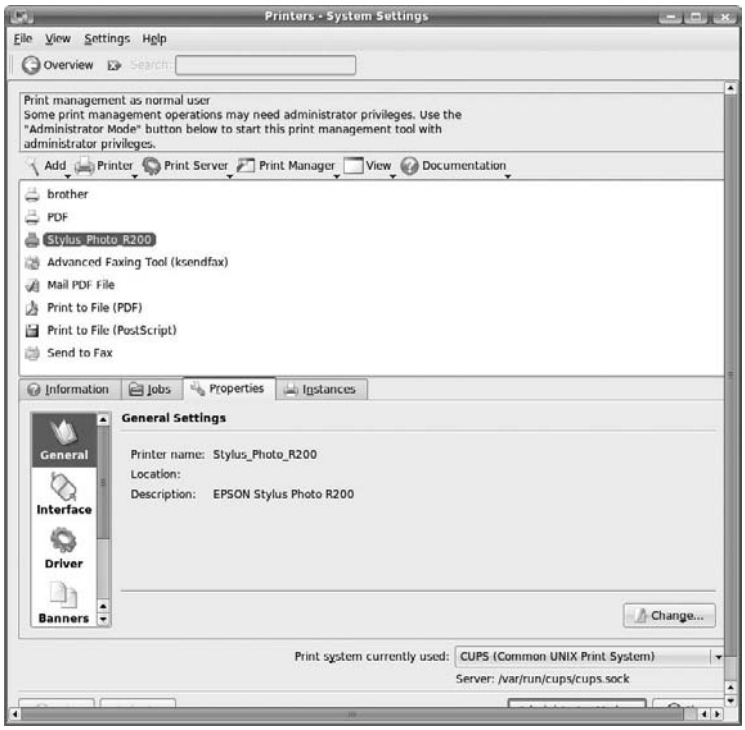


FIGURE 24-19

Configuring printer-specific settings



After you are finished reconfiguring your printer, you can click Back to return to the System Settings application, or click the Close box or select File ➪ Quit to exit the Systems Settings application altogether.

Adding Printers on KDE 4-Based Kubuntu Systems

If you are using a Kubuntu 8.10 or later system, adding a new printer or reconfiguring an existing one is done through the K Menu ➪ Applications ➪ System ➪ Printing application. Starting this application displays a screen like the one shown in Figure 24-20.

FIGURE 24-20

The KDE 4 Printer configuration utility's initial screen



To add a new printer to your system, connect the printer to your system and turn it on. KDE 4-based Kubuntu systems automatically detect most types of networked printers but do not yet automatically configure USB printers, although I'm sure that's just around the corner.

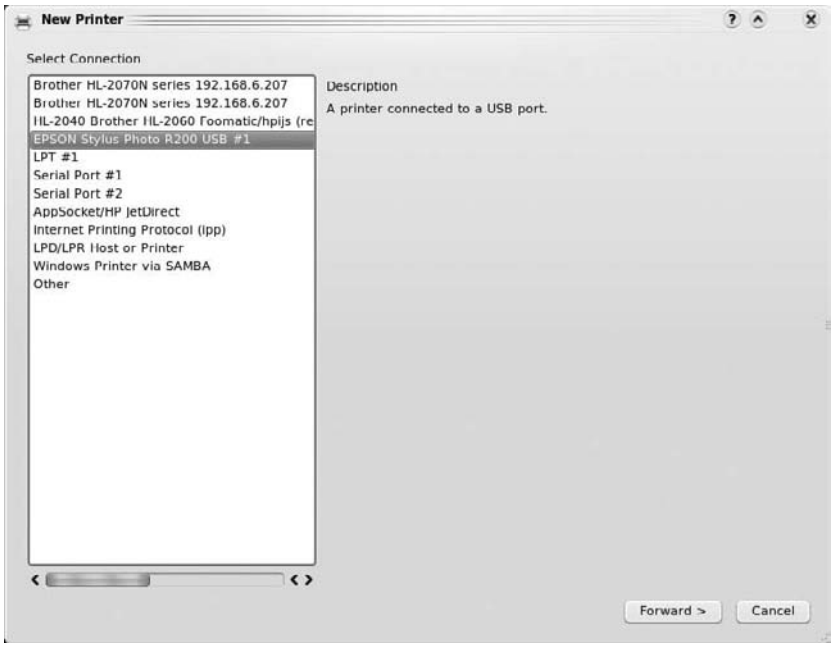
To begin configuring a new printer, select New Printer in the list shown at the left of the dialog, and then click New Printer. A dialog like the one shown in Figure 24-21 appears.

Select the option that identifies how your system will connect to the new printer:

- **USB #1:** A printer that is directly connected to a USB port on your system. This option displays only if your new printer is actually attached to a USB port on your system and is turned on. In that case, this option shows the name of the printer on that port if it could be auto-detected, as shown in Figure 24-21.
- **LPT #1:** A printer that is directly connected to your system's parallel port.
- **Serial Port #1** or **Serial Port #2:** A printer that is directly connected to one of your system's serial ports. Serial printers are extremely uncommon.

FIGURE 24-21

Identifying how you connect to your printer



- **AppSocket/HP JetDirect:** A printer that you can print directly to over the network by using a specific IP address and port. This is a common way of printing to modern standalone, network-aware printers from manufacturers such as Brother, HP, Epson, and so on. Selecting this option displays fields in which you can enter the IP address (or hostname) of the printer that you want to connect to and the port on which the printer is listening for network printing requests.
- **Internet Printing Protocol (ipp):** A printer whose queue is managed by a specific host. The Internet printing protocol (IPP) is a common alternative to network printing mechanisms such as JetDirect but supports capabilities such as authentication, access control, and host-side document formatting, and printer management. CUPS (Common UNIX Printing System), which is discussed in Bonus Chapter 12, is a specific implementation of an IPP print server that implements printing services using HTTP as a transport protocol. Selecting this option displays fields in the right portion of this dialog in which you can enter the hostname or IP address of the server to which the printer is connected, and the port on which the print server on that system is listening.
- **LPD/LPR Host or Printer:** A printer whose queue is managed by a specific host (which can be the local host) using the traditional UNIX Line Printer Damon (LPD) protocol. Selecting this option displays fields in the right portion of this dialog in which you can enter the name or IP address of the system that provides access to the printer and the name of the queue for that printer on the specified system.
- **Windows Printer via SAMBA:** A printer that is managed by a host located on a Microsoft Windows network. Such hosts are traditionally Windows servers but can also include Samba

servers running on other UNIX or Linux systems on your network. Selecting this option displays a field in the right portion of this dialog in which you can specify a URI for the printer, including the domain/workgroup, host, optional port, and printer name of the printer that you are connecting to. You can also specify whether you want to connect using a specific login and password, or whether you want to be prompted for your login and password if they are required—if authentication isn't required, you won't be prompted.

- **Other:** A networked printer that you can identify by specifying its URI (Uniform Resource Identifier).

After providing the information that is required by the connection that your printer uses, click Forward. The dialog shown in Figure 24-22 displays.

FIGURE 24-22

Selecting the manufacturer of your printer

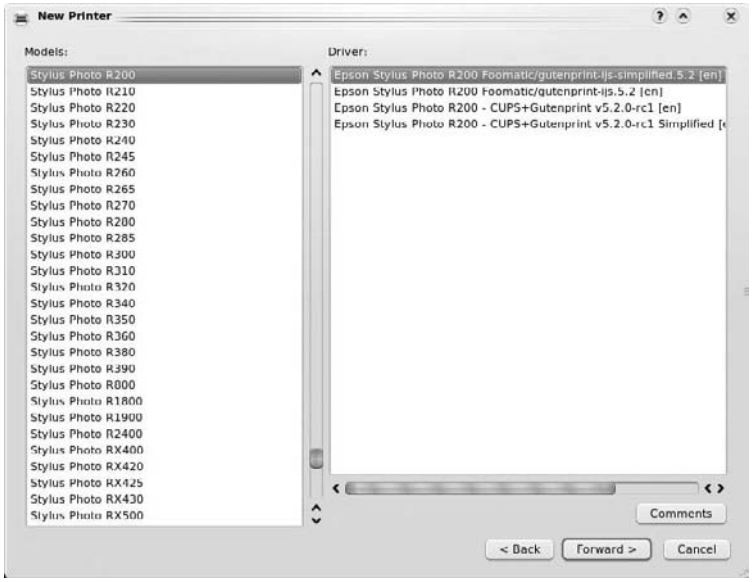


In this dialog, you select the manufacturer of your printer from a scrollable list. If you don't see an entry for the manufacturer of your printer and your printer advertises compatibility with printer languages such as HP PCL or PostScript, select GENERIC as the manufacturer of your printer.

After selecting a specific manufacturer or the GENERIC entry, click Forward. A dialog like the one shown in Figure 24-23 appears.

FIGURE 24-23

Selecting the model of your printer



Select the specific model of your printer from the Models list on the left. If the exact model of your printer isn't displayed, try selecting another model of printer from the same vendor that is as close as possible to yours. If you don't see a remotely similar model and your printer advertises compatibility with another type of printer, you may want to click Back, select the GENERIC manufacturer, and see if any of the entries on that dialog match the type of compatibility that your printer offers.

When you select a specific printer, the list at right shows the specific print drivers that are available for use based on your selection. The recommended option is preselected, and is usually the right one to try. You can always modify the printer later and use a different driver if the suggested one doesn't work optimally for you. After making your selection, click Forward. The dialog shown in Figure 24-24 appears.

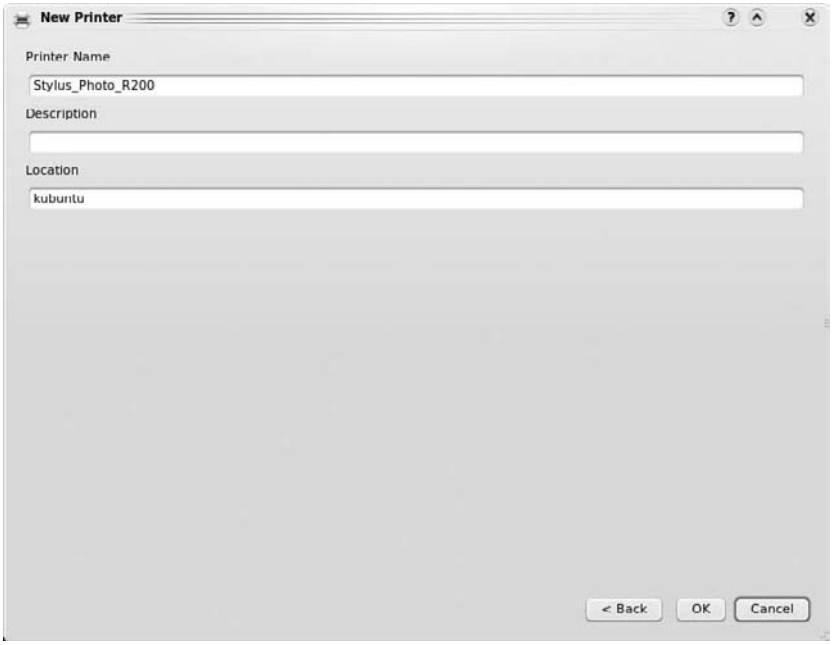
You must supply a name for your new printer because the name is used for system resources such as print queues that your system will use to manage access to the printer. A description is optional, as is a location string, which can be handy to help other users of your printer to identify a specific printer (assuming that you have more than one).

After specifying the name of your printer, optionally modifying its description, and providing an optional location, click OK to conclude the configuration process. At this point, you should be able to print to your new printer from any application on your system. To test your new printer, select its entry in the left column of the dialog shown in Figure 24-20, and click Print Test Page. A confirmation dialog displays letting you know that a test has been submitted to your printer. Click OK to close this dialog.

After you are finished defining and testing your printer, you can click the close box in the upper-right corner to close the Printer configuration application.

FIGURE 24-24

Naming your new printer



Configuring an Existing Printer on a KDE 4-Based System

At the time that this book was written, Kubuntu 8.10 and later systems offered relatively few ways of configuring a printer beyond those provided by the printer definition process. These configuration options are also available from the Printer Configuration application, which you can start by selecting the K Menu ➪ Applications ➪ System ➪ Printing. After choosing the printer that you want to modify in the list at left in the Printer Configuration dialog, your reconfiguration options from this application are the following:

- You can specify another type of printer, select another driver, or both by clicking Change beside the Make and Model label in the right portion of this dialog.
- You can make this printer the default printer by clicking Make Default.
- You can modify general printer policies, such as what happens if an error occurs and whether a banner page is printed before/after each print job, by clicking the Policies tab to display the pane shown in Figure 24-25.

After making any changes to the settings on this dialog, click Apply to save your changes, or Revert to undo your changes and redisplay the settings that you were previously using.

Many applications provide additional opportunities for printer configuration by selecting a File ➪ Print Settings or by clicking an options button on the Print dialog. For example, Figure 24-26 shows the two panes of the Options dialog for one of my printers from the OpenOffice.org Word Processor application.

FIGURE 24-25

Some basic printer configuration options on the Policies tab

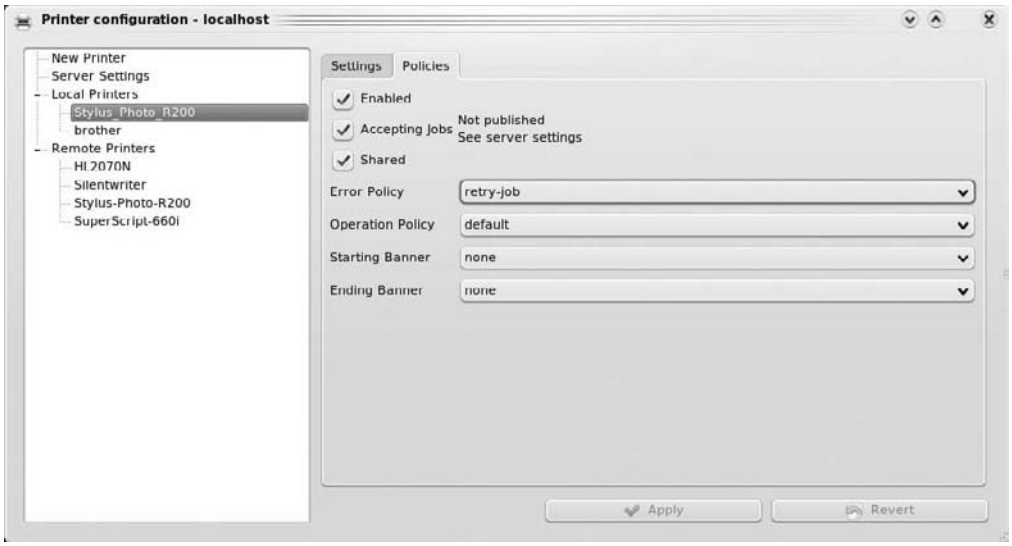
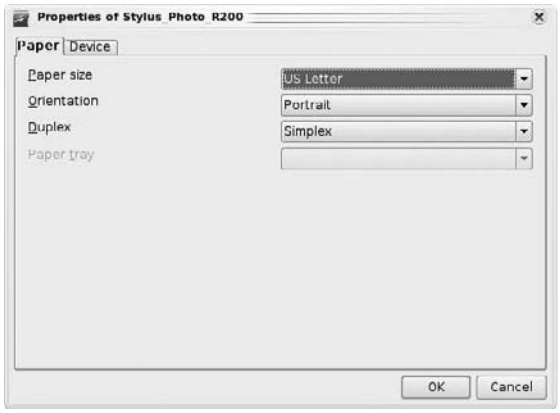
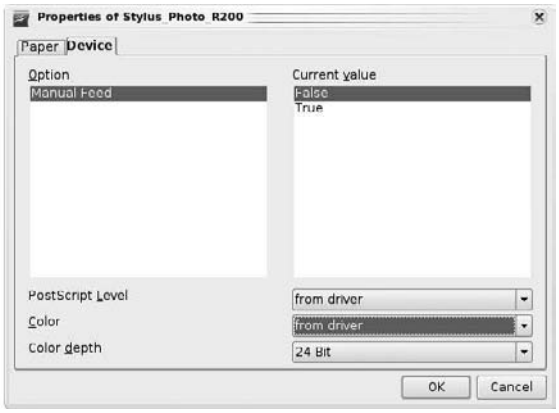


FIGURE 24-26

Setting printer options when printing from an application





Using a Scanner

Most modern scanners are USB devices and are well supported on Linux systems, as are older scanners that use a real SCSI interface. Scanners are supported on Linux systems through the SANE (Scanner Access Now Easy) interface and associated front-end and back-end applications. The home of the SANE project is `www.sane-project.org`. If you are still shopping for a scanner or are interested whether support for a specific scanner is available under Linux, see the page at `www.sane-project.org/sane-supported-devices.html`, which also provides access to an easy-to-use search engine application that provides data about specific scanners.

Ubuntu systems preinstall the XSane Image Scanner software that uses the SANE drivers but provides a graphical display of the scanning process, supports a variety of scanner configuration options, and provides both quick preview scans for alignment purposes and final scans at a variety of scan resolutions. Through the rest of this chapter, I refer to the XSane Image Scanner simply as XSane to save a few words.

KDE 3-based Kubuntu systems install Kooka (`http://kooka.kde.org`), which is an excellent scanning package, although it has been dropped from KDE 4. On KDE 4-based systems, you can install the `kdegraphics` package to get the KDE SANE libraries, but they really won't do you much good at the moment. At the time that this book was updated, your best bet on a KDE 4 system is to install XSane, which also installs the GIMP. You can then follow the process described in the "Using a Scanner on a Ubuntu System" section (coming up next) to use your scanner on your Kubuntu 8.10 or better system. Don't worry—I'm sure that someone somewhere is working on a killer scanning application for KDE 4. Let me know if you find one—I'll certainly be looking too.

For general information about SANE and scanners on Linux, see:

- **Linux Scanner HOWTO:** `http://tldp.org/HOWTO/Scanner-HOWTO/index.html`
- **SANE FAQ:** `www.xs4all.nl/~ljm/SANE-faq.html`
- **SANE Project Site:** `www.sane-project.org`

NOTE **This section is not a complete guide to every nuance of scanning software for Linux. The focus in this section is to introduce the software and common controls necessary to get a scanner working reasonably on your Ubuntu or Kubuntu system.**

Using a Scanner on a Ubuntu System

To start the XSane application, select Applications ➪ Graphics ➪ XSane Image Scanner. This displays a dialog that probes your system for available scanners and then displays a summary dialog listing all of the scanners that it has found, as shown in Figure 24-27.

To begin working with your scanner, click OK. A number of dialogs are displayed on your system. You can close the Standard options and Histogram dialogs, leaving only the main XSane dialog (see Figure 24-28) and the Preview window (see Figure 24-29) on your screen.

To get a preview scan (which is a low-resolution scan that is typically used to verify that your scanner is working, the document that you are scanning is correctly aligned, and so on), click Acquire Preview in the XSane Preview dialog. Your scanner should wake up, scan your document, and display a preview of that scan, as shown in Figure 24-29.

FIGURE 24-27

The initial XSane dialog



FIGURE 24-28

The primary XSane dialog

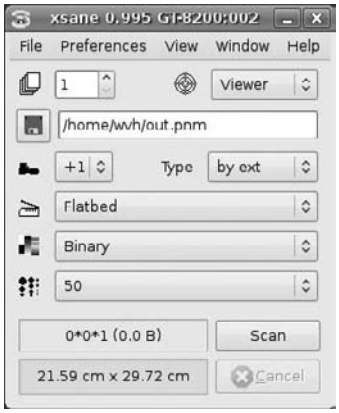
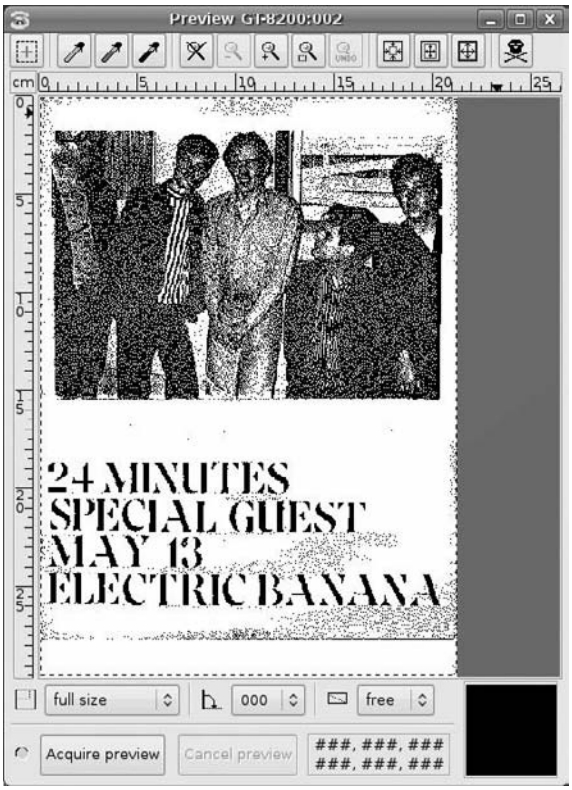


FIGURE 24-29

The XSane preview window, showing a low-res, black and white preview



Don't panic at this point—so far, I've just used XSane's default resolution and scan settings, which are quite minimal. However, if the preview is acceptable, you can proceed to do a full scan of your document. In most cases, the preview will not be of sufficient quality, so you should return to the main XSane dialog shown in Figure 24-28 to set some common scanning configuration options. The most common settings that you will want to change are:

- **Scan Mode:** Enables you to select different scanning modes by clicking on a drop-down list. The icon for this option looks like black and white, grayscale, and color gradients beside each other. Available options are binary (black and white, which is the default setting), grayscale (many shades of gray, plus black and white), and color (a full-color scan). Selecting either grayscale or color scan modes refreshes the main XSane dialog and shows additional configuration items that you can use to control brightness, contrast, colors, and so on. Figure 24-28 actually shows the main XSane dialog when a Color scan has been selected.
- **Scan Resolution:** Enables you to control the resolution at which the image is scanned by selecting a numeric value from a drop-down list. The icon for this option looks like a number of pixels of different sizes. Possible values range from 50 DPI (dots per inch), the default value, to 3200 DPI. The range of possible values depends on the scanner you are using.

The combination of the scan mode and the scan resolution determines the size of the image file that will be produced by XSane when you do a final scan. For example, a black and white scan of a U.S. Letter-sized (8.25x11 inches) document at 50 DPI produces an image file that is 13KB. A color scan at a resolution of 240 DPI increases the file size to 1.6MB. Apparently, quality does indeed come at a price.

Once you have changed your settings, preview your image again by clicking Acquire Preview in the Preview window to display a new preview that should show a substantial improvement over your initial preview. Figure 24-30 shows an improved preview.

FIGURE 24-30

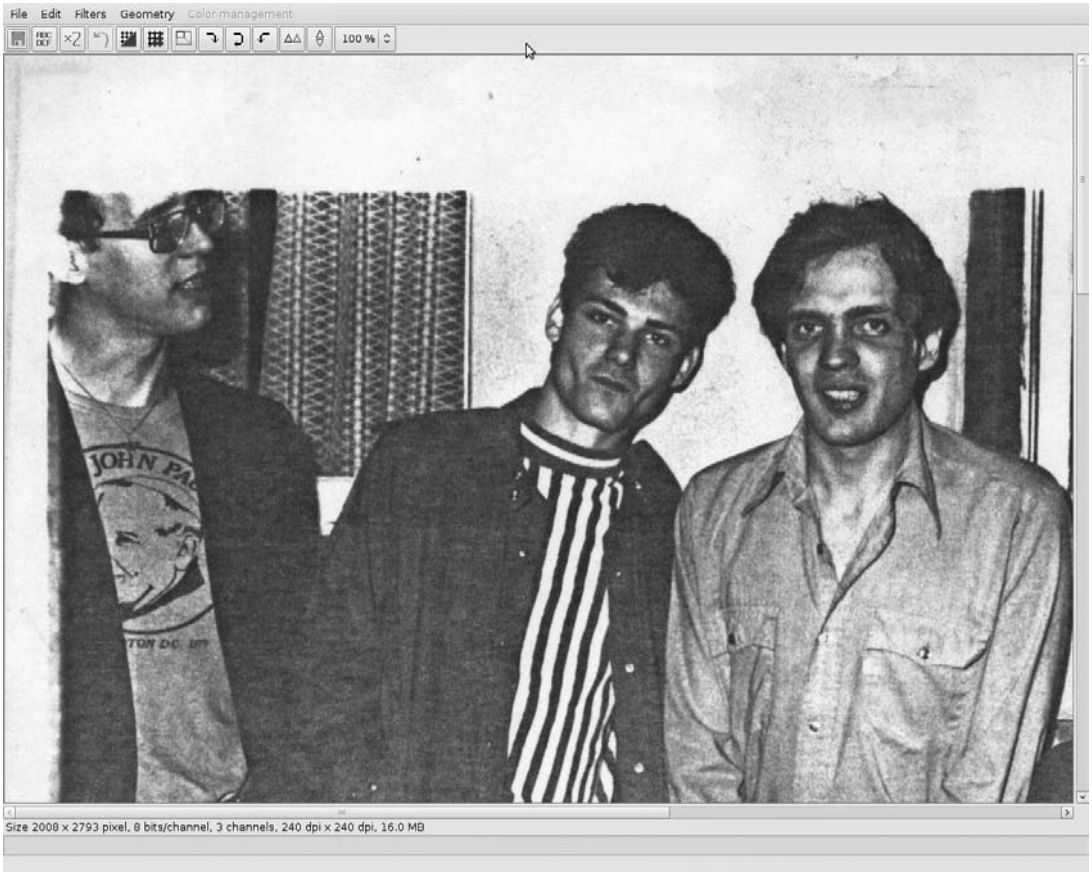
A substantially improved preview in XSane



Wow, I can now even recognize myself in this ancient poster! Assuming that the new preview looks like it will be sufficient for your purposes, close the Preview window and click Scan in the main XSane dialog. The document will be scanned using the current settings. Once the scan completes, XSane displays the scan in its standard, full-screen View window, as shown in Figure 24-31.

FIGURE 24-31

A final scan in the XSane viewer.



You may want to iterate through several scans, adjusting things like contrast and brightness by using the controls on the main XSane dialog. When you get a scan that you are happy with, you can save the scanned image by clicking File ➪ View, which displays a dialog that enables you to specify a filename and also the image format in which the scan will be saved.

TIP **If your Ubuntu system is having a problem locating your scanner or you need to know the name of the device node associated with your scanner, use your favorite package manager to install the `sane-utils` package, which provides a convenient utility called `sane-find-scanner` that does an excellent job of finding SCSI and USB scanners. Installing the package also installs man pages for `sane-scsi` and `sane-usb`, which provide a great selection of useful suggestions for resolving problems locating SCSI or USB scanners.**

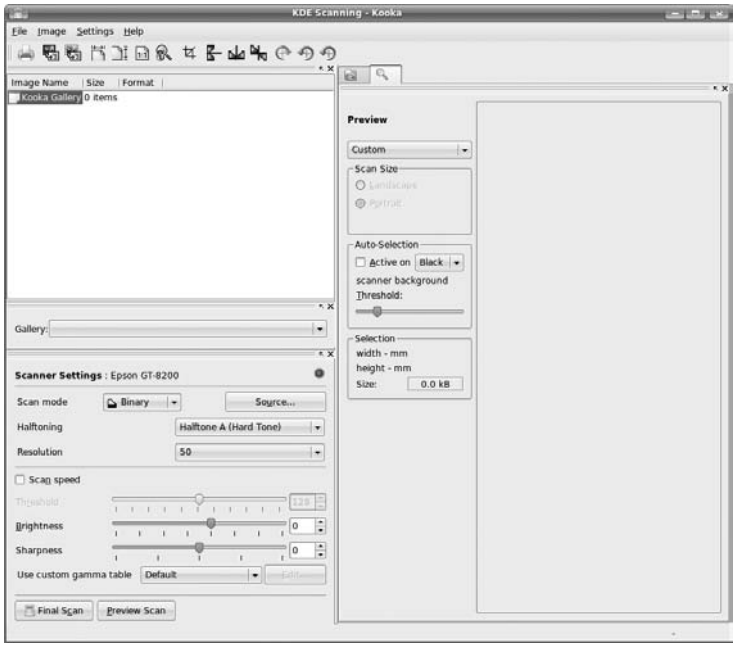
Using a Scanner on a KDE 3-Based System

Kooka (`http://kooka.kde.org`) is an excellent scanning package for KDE 3-based systems that unfortunately hasn't made the cut for KDE 4. As sad as that is, it doesn't diminish Kooka's usability and power on KDE 3-based systems such as Kubuntu 8.04 and earlier.

To start Kooka, select the K Menu ➪ Graphics ➪ Kooka - Scan & OCR Program menu item. This displays a dialog that probes your system for available scanners and then displays a summary dialog listing all of the scanners that it has found. It also provides an optional checkbox that tells Kooka to always use the scanner that has been detected, which enables you to skip this dialog in the future and can be useful if you only have one scanner. Click OK to proceed; he initial Kooka screen appears, as shown in Figure 24-32.

FIGURE 24-32

The primary Kooka dialog

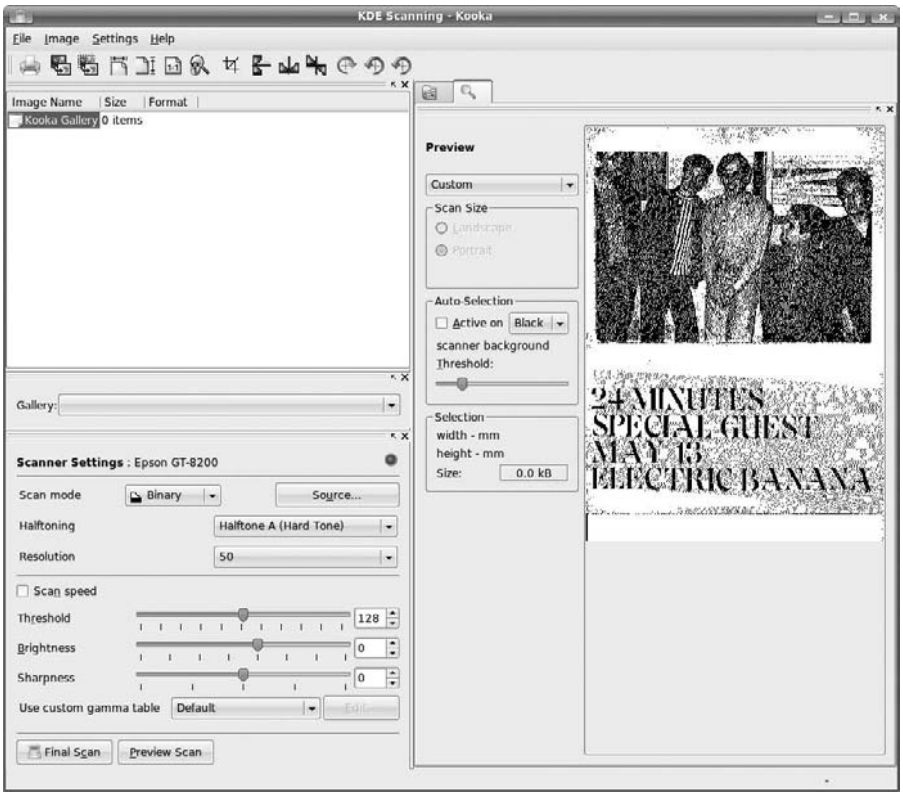


To get a preview scan (a low-resolution scan that is typically used to verify that your scanner is working, the document that you are scanning is correctly aligned, and so on), click the Preview tab in the right pane (its icon is a magnifying glass), and then click Preview Scan at the bottom left in the main Kooka dialog. Your scanner should wake up, scan your document, and display a preview of that scan, as shown in Figure 24-33.

TIP **On some of my Kubuntu systems, Kooka has the irritating habit of using a default window size that is too small to display the Final Scan and Preview Scan buttons in the lower-left corner. It's tricky (that is, impossible) to use Kooka without having access to these buttons. If you don't see them, resize your window so that they are visible. Voilà—usable software!**

FIGURE 24-33

Kooka showing a low-res, black and white preview



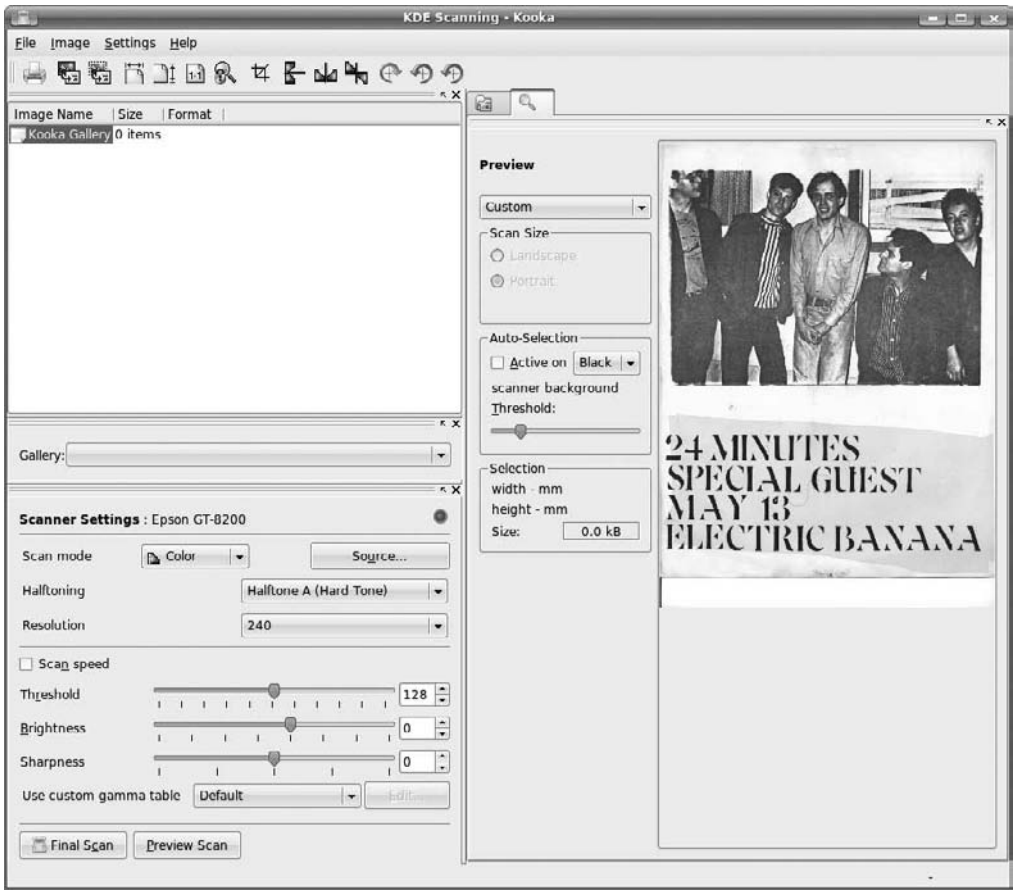
So far, I've only used Kooka's default resolution and scan settings, which are minimal. If the preview is acceptable, proceed to do a full scan of your document. In most cases, though, the preview will not be of sufficient quality, so you'll need to set some common scanning configuration options, which are located in the Scanner Settings portion of the Kooka window. The most common settings that you will want to change are:

- **Scan Mode:** Select different scanning modes by clicking on a drop-down list. Available options are binary (black and white, which is the default setting), grayscale (many shades of gray, plus black and white), and color (a full-color scan).
- **Resolution:** Control the resolution at which the image is scanned by selecting a numeric value from a drop-down list. Possible values range from the default 50 DPI (dots per inch) to 3200 DPI. The range of possible values depends on the scanner you are using.

Once you have changed your settings, preview your image again by clicking Preview Scan again to display a new preview that should show a substantial improvement over the initial one. Figure 24-34 shows an improved preview.

FIGURE 24-34

A substantially improved preview in Kooka



Assuming that the new preview looks as though it will suffice or your purposes, click Final Scan. The document will be scanned using the current settings.

Once the scan completes, Kooka displays the dialog shown in Figure 24-35.

The dialog enables you to specify the image format in which you want to save the scanner image. Kooka supports most standard lossless graphics and scan formats, including BMP, PNG, PPM, and TIFF, and also supports popular lossy formats such as JPEG. Selecting a format and clicking OK saves your new scan with a default name, which will be listed in the Kooka Gallery pane at the top left of the Kooka window. You can then click the File/Folder tab above the right pane in the Kooka window to display the high-quality scan, as shown in Figure 24-36.

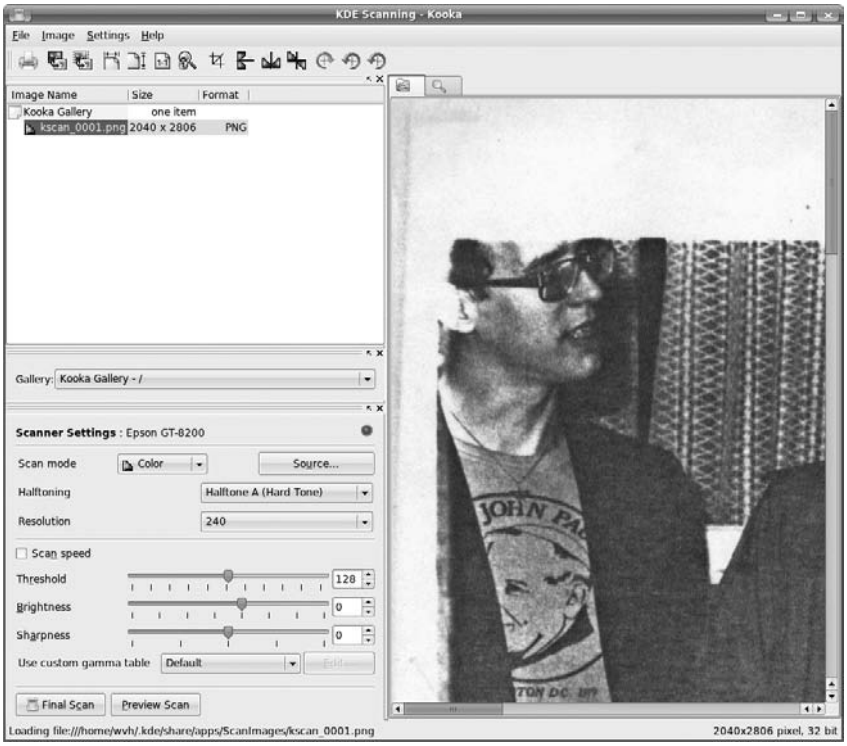
FIGURE 24-35

Selecting the graphics format in which to save your scanned image



FIGURE 24-36

A final scan in Kooka



You may want to iterate through several scans, adjusting things like Threshold, Brightness, and Sharpness, using the controls in the Scanner Settings portion of the Kooka dialog. Remember to delete any scans that you don't actually want—scanned documents can consume a significant amount of disk space, especially at higher resolutions.

Adding Internal Disks and CD/DVD Drives

Portable USB and FireWire hard drives are common today, as are USB flash drives (which look like small hard drives to your system). These portable, removable storage devices simplify backups, transporting large amounts of data, and transferring large amounts of data from one system to another. For the most part, these removable devices are correctly identified by Ubuntu and Kubuntu Linux systems, which will either mount new devices or prompt you for the action to take, as described in the portions of Chapter 21 on digital device detection.

While portable, external devices are convenient, you may eventually want to add more storage to your system permanently, which usually means adding new drives to free connections that are available on your system. These interfaces are faster than the interfaces used for external storage, and also keep the number of cables running to and from the exterior of your system to a minimum.

This section explains how to add new internal drives to your systems. The most common types of hard drives found in desktop and laptop PC systems today are EIDE/ATA (Enhanced Integrated Drive Electronics/Advanced Technology Attachment) and SATA (Serial Advanced Technology Attachment) drives. Because of the popularity of SATA drives, EIDE/ATA drives are also often known as PATA (Parallel Advanced Technology Attachment) drives, which highlights the basic difference between these two types of drives (parallel versus serial) and also delivers a nicely parallel acronym at the same time. SCSI (Small Computer System Interface) hard drives are another possibility, but are much less common in desktop systems, although they are still often used in server systems. Most CD and DVD drives are EIDE/ATA drives, using a standard referred to as ATAPI, although SCSI and SATA CD and DVD drives are also available.

To add any of these drives internally to your system, you will have to shut down the system and remove its side cover to give you access to the motherboard and to the portions of the case where hard drives can be mounted, which are usually referred to as *drive bays*.

The next few sections explain how to add extra drives of these various types to your system, describe how to identify the device nodes associated with those drives when you reboot your system, and provide an overview of partitioning and creating filesystems on these devices.

Adding EIDE/ATA Drives

Once your system is turned off and open, locate the IDE cables that are attached to the EIDE interfaces on your motherboard (or, in rare cases, to an IDE or EIDE controller card). IDE cables are flat, 40-pin ribbon cables with two flat, notched connectors near one end and a similar connector at the other end. Most PCs have two IDE interfaces, known as primary (IDE-1) and secondary (IDE-2). The cable from each IDE interface can be attached to a maximum of two hard or CD/DVD drives. If two drives are attached to a single cable, they must be configured as master and slave drives by connecting pins known as *jumpers* that are located on the back or bottom of the hard drive or CD-ROM drive. *Master* is the term used for the first drive on an IDE interface; *slave* is the term used for the second drive on the IDE interface.

To add a new drive to your system, you must set the jumper (or jumpers) on the new drive so that it does not conflict with any existing drive on a given IDE interface in your system. To do so, you must

know how the current drives are connected. You can determine how the existing hard drive(s) and CD/DVD drive(s) are connected in one of two ways: automatically (by checking the BIOS settings on your system) or manually (by simply checking the connections and jumpers on the existing drives). Because you've (hopefully) already turned the PC off and the ways of collecting information from a PC BIOS differ depending on which BIOS your system uses, we'll focus on looking just at the cables and jumpers.

By tracing cables and examining jumpers, you can manually identify how the drives are connected. To do so, follow each of the IDE cables from your motherboard back to any hard drives or CD/DVD drives connected to them. The hard drive that your PC boots from is known as the *boot drive*, and it is usually connected as the master on the first IDE interface on your system. If your CD/DVD drive is connected to the same IDE cables as your hard drive, it is probably jumpered as the slave on the primary IDE interface (IDE-1). If your CD-ROM drive is connected to your other IDE cable (IDE-2), it is probably jumpered as the master on that interface.

TIP **If there are already two devices attached to an IDE interface and your PC was working correctly before, skip that IDE interface—it's already fully populated and correctly configured.**

Once you know how your current drives are connected, you can proceed to jumper your new drive so that it does not conflict with existing drives in your system. A guide to the jumpers on your hard drive that will cause the drive to be recognized as a master or slave drive is usually printed on the top of the hard drive.

CAUTION **Be careful when reading the jumper information; although it is usually written from left to right, just as the jumpers are physically located on the hard drive, some jumper settings are listed relative to a jumper key, which is a single pin on the left or right side of the hard drive jumpers.**

After you jumper your new drive, mount it in a drive bay, attach it to the selected IDE cable in your system, and provide power to the drive by connecting it to one of your system's power feeds. At this point, you can power up the system, boot into Linux, and then follow the instructions in the section later in this chapter entitled "Locating, Partitioning, and Formatting New Drives." Some common problems at this point are the following:

- If you have trouble booting your system after adding a new drive, see the section entitled "Troubleshooting Boot Problems after Adding New Drives" later in this chapter.
- The most common source of problems when adding a new EIDE/ATA drive to an existing system is jumpering it incorrectly. If a new drive is jumpered to be the same (master or slave) as an existing drive, your system will usually not be able to see either drive. You can verify that the new drive is jumpered, attached, and detected correctly by entering the BIOS immediately after rebooting the PC and examining the primary and secondary IDE interfaces.
- If you have problems finding the new drive after booting, reboot your system and press a key such as the F2 or Delete key to enter your system's BIOS setup screens. (A message that identifies the key that is required to access your system's BIOS is usually displayed when you boot your computer.) The Main or System section of these screens should enable you to probe your IDE or SATA drives to verify that your system can actually find them. If your system doesn't see the drive(s) at all, check the power and IDE/SATA connections to the drive.

You will not be able to use any new drive on your system until you have partitioned and formatted the drive. Some drives come preformatted with an NTFS filesystem, which is probably not a filesystem that you want to use on an internal drive on a Linux system. See the section later in this chapter entitled "Locating, Partitioning, and Formatting New Drives" for detailed information about partitioning and formatting your new drives.

Hard Drive Names on Linux Systems

The traditional IDE subsystem on Linux systems recognized EIDE/ATA hard drives as devices with names of the form `/dev/hdX`, where *X* was a letter that reflected the sequence in which drives were located on your system's IDE interfaces. SCSI hard drives used their driver subsystem, using names of the form `/dev/sdX`. SATA also used the SCSI subsystem and, therefore, followed the same naming convention as most removable Flash and USB drives. However, in 2007, the existing IDE subsystem in the kernel was replaced by a newer PATA subsystem, which uses the `/dev/sdX`. naming convention. This is a superior subsystem, although it certainly confused a number of people when drives formerly known as `/dev/hdX`. started showing up with names in the form `/dev/sdX`. after a kernel upgrade, especially if their GRUB configuration file wasn't updated at the same time. This should all be water under the bridge now, but it's worth explaining because hard drive naming conventions often confuse people in the first place, let alone after drives "rename themselves."

Adding SATA Drives

SATA drives are much easier to attach and use than EIDE/ATA or SCSI drives, primarily because no jumpering is required. SATA drives are automatically identified by your system based on the number of the SATA port that they are attached to. Most motherboards with SATA support have either two or four SATA ports that feature the small L-shaped SATA connector in their center. Standalone SATA controllers are also available that typically provide four SATA ports but can provide many more.

SATA has many advantages over traditional EIDE/ATA drives (which I'll call PATA drives in the rest of this section). First and foremost, each SATA drive has a separate cable connecting it to the controller, reducing the shared cable, master/slave model used by PATA drives. Although performance is approximately the same between SATA II and PATA drives, the dedicated connection and controller for each SATA drive can provide higher overall throughput in disk-intensive environments. SATA cables can be up to 39 inches long, while PATA cables are supposed to be limited to 19 inches. The fact that SATA cables are much smaller than PATA ribbon cables makes them easier to route aesthetically within the tight confines of many computer cases. And finally, the biggest advantage of SATA over PATA drives is that SATA drives can be hot-swapped, which makes them eminently suitable for use in RAID or other environments with 24x7 availability requirements.

To add another SATA drive to your system, simply mount it in an available drive bay, connect it to an available SATA port, and provide power to the drive by connecting either a traditional power connector or a special SATA power connector, depending on the drive that you're using. At this point, you can power up the system, boot into Linux, and then follow the instructions in the section later in this chapter entitled "Locating, Partitioning, and Formatting New Drives." If you have problems booting your system after adding a new drive, see the section entitled "Troubleshooting Boot Problems after Adding New Drives" later in this chapter.

TIP **Some motherboards provide two different sets of SATA ports, one set for hardware RAID support and another set of normal SATA ports. Do not use the hardware RAID SATA ports unless you're sure that the hardware RAID controller is well supported under Linux and you're sure that you want to use hardware RAID. I usually prefer Linux software RAID to Linux hardware RAID because hardware RAID disk layouts are controller-specific. Using hardware RAID often means that if your motherboard fails, you will have to replace it with an identical motherboard or another with the same hardware RAID controller. In most cases, you will not be able to use your existing RAID drive(s) with a different hardware RAID controller.**

Adding SCSI Drives

Alas, poor SCSI, I knew thee well. SCSI drives have long been the traditional drive of choice for RAID boxes or in systems with 24x7 availability requirements. The primary reasons for this are high performance (SCSI drives are still faster than current SATA II drives), SCSI's traditional support for hot-swapping, and higher quality control (some manufacturers even offer longer warranties for SCSI drives for this reason). The downside is that SCSI drives are substantially more expensive than today's SATA II drives, which offer the same hot-swap capabilities. Some people would rather fight than switch—your call. Traditionally, only very high-end motherboards feature integrated SCSI controllers—SCSI controllers in most systems are standalone SCSI controller cards.

SCSI connectors and controllers vary according to the version of the SCSI standard that they conform to and the number of drives that they can support. Various types of SCSI connectors have evolved over the years, which are roughly associated with different versions of the SCSI standard and are specifically associated with different performance characteristics. The most common of connectors are Centronics-style parallel and IDC-50 connectors (typically associated with 8-bit SCSI-1 drives), DB-50 connectors (typically associated with 8-bit Fast SCSI-2 drives), and DB-68 connectors or SCA-80 (typically associated with 16-bit, Fast/Wide SCSI-3). SCSI-1 and SCSI-2 controllers support up to seven drives per controller, while SCSI-3 and later SCSI controllers support up to 15 drives per controller.

Each SCSI drive in a system must be uniquely identified by setting it to a unique ID number. In most cases, you must set one or more jumpers on the drive to indicate its ID, but if you have more money than I do and are using a multiport SCSI controller, that controller may identify each drive based on the port that it is connected to. The last SCSI drive in a chain must also be terminated (that is, drives that are installed sequentially, with each cabled to the next), although some drives and SCSI controllers do this automatically.

To add another SCSI drive to your system, set it to a unique ID if necessary, mount it in an available drive bay, connect it to your SCSI controller or to the next SCSI drive in a chain, and provide power to the drive by connecting a traditional power connector. At this point, you can power up the system, at which point you should see the BIOS on your SCSI controller identify the IDs on all connected SCSI drives. Once you boot into Linux, you can follow the instructions in the section later in this chapter entitled "Locating, Partitioning, and Formatting New Drives." If you have problems booting your system after adding a new drive, see the next section.

Troubleshooting Boot Problems after Adding New Drives

It is common that adding a new hard drive to your system will prevent it from booting immediately. Much of this has to do with settings in your computer's BIOS or other boot environment that enumerate the hard drives that are available and tell it in which order to try various disks and interfaces in order to boot. The next few sections discuss a few common problems and how to resolve them.

Changes to BIOS Hard Disk and Boot Device Priority Settings

Most BIOS menus feature a "Boot" or "Boot Settings" menu or screen that contains entries that lead to configuration screens for Hard Disk Drives and Boot Device Priority. The Hard Disk Drives screen usually enumerates all of the hard drives in your system. Frequently, the first hard drive listed is the first that can be identified on the Boot Device Priority screen. Make sure that the first hard drive in this list is the drive that actually contains your operating system. Next, make sure that the hard drive that contains your operating system is actually listed in the Boot Device Priority screen, and that the entries on this screen appear in the right order. Many Boot Device Priority screens only provide three or four slots for possible boot devices, and I usually squander the first two of these for "Floppy Disk" and "CD-ROM" boot devices. Whatever boot device order you prefer, make sure that the hard drive that contains your operating system is present in this list and appears in the right sequence.

Problems after Adding IDE Drives to a SATA System

Most systems are configured to try to boot from IDE drives before SATA disks. If you have added IDE disks to a system that formerly only contained SATA drives, make sure that the sequence in which your system tries to boot from different disks is correct. I had a humorous experience a while ago when recycling some old IDE disks for use on one of my SATA systems, because I hadn't wiped the IDE disks. I'd forgotten that one of them contained an older Linux installation and was quite surprised when my Ubuntu system suddenly booted SUSE.

Problems Related to SATA Drive Name Changes

Adding SATA drives to an existing system can temporarily cause boot problems for various reasons, the most amusing of which is because of how SATA drives are detected and named. SATA disks not only use SCSI emulation (and thus have names beginning with `/dev/sd`), but the letter that follows `/dev/sd` in their names is assigned based on the sequence number of the specific SATA port that they are attached to. If whoever built your system was careful, your primary SATA disk will be attached to SATA port 1, and thus will always be assigned the name `/dev/sda`. If whoever built your system wasn't so careful and attached your boot drive to SATA port 4, your system will boot fine with a single drive, but adding another SATA drive will cause that drive to be assigned the name `/dev/sda`, and your existing SATA drive will be called `/dev/sdb`. Adding more than one new SATA drive will continue to demote the name of your original SATA disk. This generally confuses both the BIOS, as described in the previous paragraph, and certainly the boot loader used by Ubuntu systems, GRUB.

If you encounter this problem, there are two possible solutions:

- Open up your system again and connect your original drive to the first SATA port on your motherboard. This will guarantee that it will always be named `/dev/sda`. You may still have to change the BIOS settings, but you won't have to modify anything in your drive's boot configuration files.
- Change the boot settings used by the Linux GRUB boot loader to reflect the new name of your drive.

The first option is almost always preferable unless you have a very good reason for simply modifying the GRUB boot settings. If you insist on doing the second, and modifying GRUB, see the next section.

Temporarily Changing GRUB Boot Loader Information

If you need to change the GRUB boot loader's idea of the root partition on your system, first correct the boot order of the disks in your BIOS as described in the section earlier in the chapter entitled "Changes to BIOS Hard Disk and Boot Device Priority Settings." This will enable your system to begin to boot from your existing hard drive, which will almost certainly fail once GRUB tries to mount its root filesystem, which is probably identified with a name such as `/dev/sda1`; this is still the right partition but is now located on a disk with a different name. To temporarily modify GRUB's boot settings to work around this problem, you must do the following:

1. Press Escape during the boot process to display the GRUB menu.
2. Press E to display the boot commands associated with your default boot entry.
3. Use the arrow keys to highlight the line beginning with the word "kernel."
4. Press E to edit that entry, and use the arrow keys and the Backspace or Delete key to change the name of the partition identified as the root partition (through the `root=` keyword). For example, if your disk was originally `/dev/sda` but is now `/dev/sdd`, the root partition is probably still identified as `/dev/sda1`. Change this to `/dev/sdd1` and press Return to return to the boot options screen.

TIP **If your system already had multiple disks and you were normally mounting partitions from all of them, your system may not come up in multiuser mode correctly because any partition names in your `/etc/fstab` file will be wrong. You may also want to temporarily add the word `single` to the end of the GRUB kernel entry to bring up the system in single-user mode so that you can correct the names of the partitions in `/etc/fstab`.**

5. Press B to boot using your modified GRUB boot options.

Once your system comes up correctly, you can use a text editor to permanently modify the file `/boot/grub/menu.lst`, which is GRUB's configuration file. You should not only modify the boot settings for your default GRUB boot selection but also modify the template values in this file that are used to prepare a new entry each time a new kernel is installed on your system.

You should also check the contents of the `/etc/fstab` file on your system and make sure that it reflects the current disk names in use on your system. It is not uncommon for a system to come up without swap space and identify its root partition incorrectly (in commands such as `df`) because entries in the `/etc/fstab` file are incorrect. During the boot process, the root partition is already mounted by the time the `/etc/fstab` file is used to mount other disks.

Problems after Adding PATA/SATA Drives to a SCSI System

The last common problem that you may encounter occurs when adding SATA or PATA drives to a system that boots from SCSI disks. Both PATA and SATA drives now use the same naming convention as SCSI disks. PATA and SATA disks are also named before SCSI disks during the kernel boot process. The combination of these two facts means that adding one or more PATA/SATA disks will cause your SCSI disk to be renamed. For example, if your system was configured to boot from the real SCSI disk, `/dev/sda`, and you add two SATA drives, the name of your SCSI disk will now be `/dev/sdc`. This leads to the same sort of BIOS and GRUB boot loader problems discussed earlier in the "Problems Related to SATA Drive Name Changes" section, but you will have to solve them using the solution described in the section "Changing GRUB Boot Loader Information." See that section for an explanation of how to modify your GRUB settings and resolve this sort of problem.

Locating, Partitioning, and Formatting New Drives

After rebooting your system to locate new internal drives, you'll probably want to use the storage that those new disks provide. This requires that you identify the Linux device nodes associated with them, and then use various software packages to partition them and format those partitions using your filesystem of choice.

As you might expect, there are some great command-line utilities for displaying information about the disks that are attached to your system, partitioning them, and formatting those partitions so that they contain a filesystem that you can mount in order to actually use the storage that they offer. Unfortunately, the story is not so good for graphical utilities that enable you to do the same sorts of things. Earlier releases of Ubuntu featured a graphical Disks utility on the System ➪ Administration menu (which you could execute from the command line as `disk-admin`), but it is no longer provided on Ubuntu systems, and its most likely replacements, Disk Manager (`http://flomertens.free.fr/disk-manager/index.html` or `https://launchpad.net/disk-manager`, for the Ubuntu take on it) or GNOME Format (`http://live.gnome.org/gnome-format`) are not part of the Ubuntu distribution yet. The story is better on KDE 3-based systems such as Kubuntu 8.04 and earlier, which provide support for formatting and partitioning disks through the Advanced tab's Disk & Filesystems applet. Life gets hard again for KDE 4–based systems such as Kubuntu 8.10 and later, because this applet is not yet available for the KDE 4 System Settings application. Therefore, I've taken the somewhat heretical (for Kubuntu users)

approach of focusing on the GNOME Partition Editor (`gparted`) as the "right" graphical tool for both Ubuntu and Kubuntu desktop systems. (Server users without a GUI will have to use the command-line tools discussed in this section anyway.) I apologize in advance to any KDE fanatics who resent violating the ethnic purity of their KDE desktop system with a GNOME package, but I think it's smarter to use the best package for the job, regardless of its graphical roots. The GNOME and KDE wars are as meaningful as the "Tastes Great/Less Filling" conflict that rocked the world of bad American beer in the late twentieth century. Installing this tool on a Kubuntu system does drag along the standard amount of GNOME baggage (including, surprisingly, GNOME's graphical `gksu` application for requesting and entering your password), but the disk space that you'll sacrifice is well worth it in terms of the usability and power of the GNOME Partition Editor. The GNOME Partition Editor is just a GNOME front end to the command-line `parted` tool, but having a GUI for `parted` turns a somewhat arcane command-line experience into a task that anyone who knows enough to add a new disk to his system should have absolutely no problem with.

TIP **A QT front end to the `parted` tool, cleverly named `qtparted`, is available in the Ubuntu repositories, but it doesn't seem to have been updated for quite a while and doesn't really work very well (when it doesn't simply hang when performing some operation). I'd avoid this package.**

The next three sections explain how to use the GNOME Partition Editor and command-line utilities to figure out what drives are attached to your system, the partitions they contain, how to partition or repartition those drives, and how to format those partitions.

NOTE **Chapter 4, "Basic Linux System Concepts," provides an overview of Linux system concepts such as disk partitions, disk geometry, and related terms such as sectors and cylinders, filesystems, filesystem mount points, and so on. If you encounter unfamiliar terms in this section, refer to Chapter 4 for more information.**

Using the GNOME Partition Editor to Examine and Format Drives

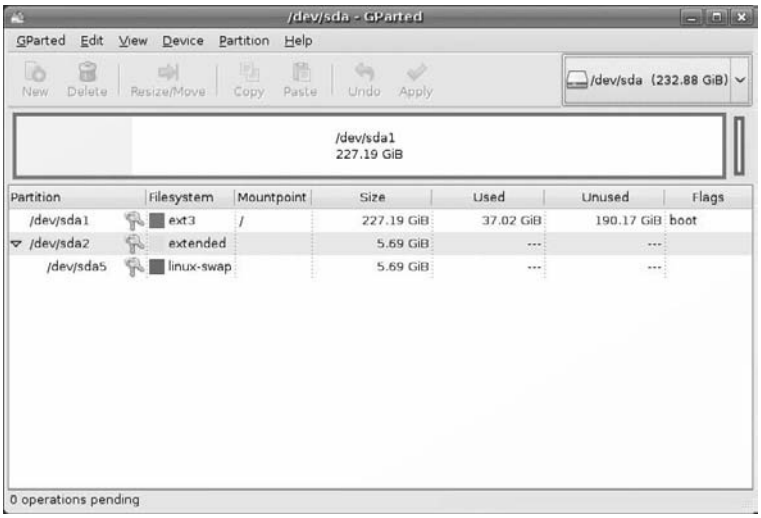
As mentioned in the introduction to this section, the GNOME Partition Editor is a graphical front end to the command-line `parted` tool, and is a great tool for use on Ubuntu and Kubuntu systems of any type. This tool is not installed by default, but you can resolve that shortcoming by using your package management tool of choice, discussed in Chapter 23, to install the `gparted` package. Once installed, the GNOME Partition Editor provides a graphical view of all of the storage devices that are currently installed on your system and makes it easy for you to determine basic information about their capacity, any partitions that they contain, and the Linux device nodes associated with the device and each partition. It also enables you to add, remove, and format hard disk partitions.

On a Ubuntu system, you can start the GNOME Partition Editor by selecting System ➪ Administration ➪ Partition Editor. On a KDE 3-based Kubuntu system, you can start the partition editor by selecting the K Menu ➪ System ➪ Partition Editor menu command. On a KDE 4-based system, you can start the partition editor by selecting the K Menu ➪ Applications ➪ System ➪ Partition Editor menu command. After entering your administrative password, the dialog shown in Figure 24-37 appears.

The panes display information about the storage device that is currently selected. The drop-down list in the upper-right corner lists all of the storage devices that are currently detected on your system, and provides information about the storage capacity of each device and the Linux device node through which that device can be accessed. Once a device is selected from the drop-down list, the pane at the bottom displays a summary of how that storage device is currently partitioned, as shown in Figure 24-37. The top pane shows a graphical view of the partitions on the selected disk, while the bottom pane identifies each partition, the device node through which it can be accessed, its format, and so on. For example, the partition list for the disk that is selected in Figure 24-37 shows that it is the boot disk for the current system, because one of its partitions is mounted as `/` and, therefore, contains the current system's root filesystem.

FIGURE 24-37

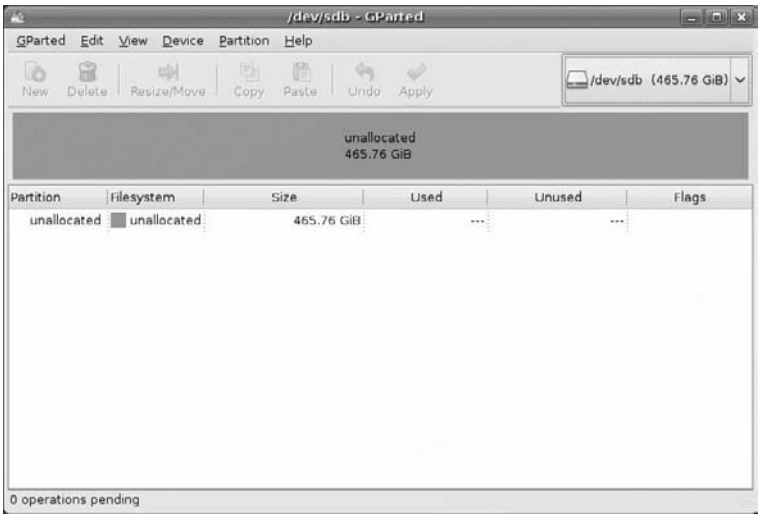
Device information in the GNOME Partition Editor



After adding a new disk to your system, select each device in the drop-down list to examine its current partitioning until you find the device that is associated with the new drive. Figure 24-38 shows the entry for a new disk that has not been partitioned.

FIGURE 24-38

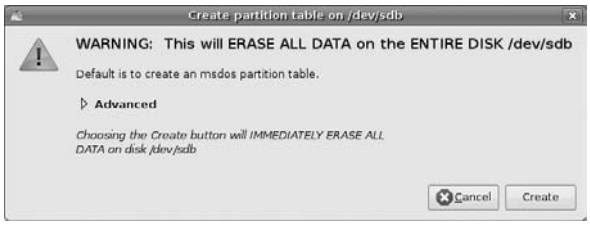
Information about a new, unformatted disk



When working with a new, unformatted disk, you will first need to create a partition table so that you can subsequently create one or more partitions on the disk. To do this, select Device ➪ Create Partition Table. A dialog appears, as shown in Figure 24-39, warning you that creating a partition table on the device will erase all data on the disk (because doing so deletes any existing partitioning information that may currently be present on this disk).

FIGURE 24-39

The warning dialog that appears when creating a partition table



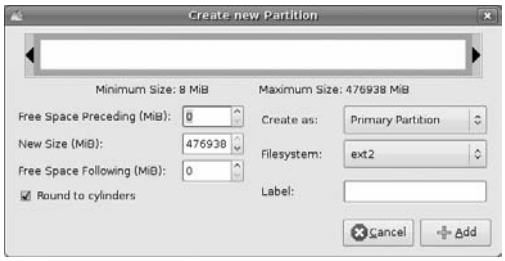
After double- (and triple-) checking that you are indeed working with the correct disk, click Create to create a partition table on the device. Once the partition table has been created, a screen like the one shown in Figure 24-38 reappears. You're now ready to begin creating partitions on your new disk.

NOTE **By default, Linux systems use a partition table that is in the traditional MS-DOS format, which makes it possible to use disks that contain both Linux and MS-DOS/Windows partitions in a computer system that boots both Linux and Microsoft Windows. You can also create partition tables in different formats in the GNOME partition editor, which enable these disks to be used in systems that can boot both Linux and other operating systems, such as Solaris, AIX, and so on. To select one of these other partition table formats, click the arrow beside the Advanced entry on the warning dialog shown in Figure 24-39, and select one of the other partition table formats. Using a partition table format other than MS-DOS is extremely uncommon, but the option is available if you need it.**

To create a partition on the current disk, first select the entry for the unallocated space on that disk in either the graphical or list portion of the dialog shown in Figure 24-38. Next, either right-click and select New from the context-sensitive menu that displays, or select the Partition ➪ New menu item to display the dialog that enables you to specify the characteristics and size of that partition, as shown in Figure 24-40.

FIGURE 24-40

Creating a partition



The dialog enables you to specify several different types of information about the partition:

- **Free Space Preceding:** Enables you to leave some free space on the disk before the partition that you are creating. Wanting to do this is uncommon, but it can be useful if you are formatting a disk on which you subsequently want to install Microsoft Windows, which generally wants to be the first operating system on the disk (the big bully). Leaving free space before your new partition will leave unallocated space on the disk, which Windows can subsequently format and use.
- **New Size:** The size of the partition that you are creating. By default, the dialog assigns all available free space to any new partition. If you want to create multiple partitions on your disk, you need to reduce the number that is initially displayed to leave unallocated space in which you can create one or more other partitions. If you reduce the size of the partition that you are creating so that it does not use all of the unallocated space on the disk, the graphical display at the top of the dialog changes to reflect the size of the partition that you are creating and the amount of unallocated space that remains on the disk.
- **Free Space Following:** Enables you to leave some free space on the disk after the partition that you are creating. The number in this field is adjusted automatically if you do not allocate all of the available space in the unallocated space that you selected before beginning the process of creating a new partition.
- **Round to Cylinders:** Selected by default, this option causes the size of new partitions to be rounded to an appropriate number of disk cylinders to enable Linux to use your disk as efficiently as possible without wasting space.
- **Create as:** Enables you to specify the type of partition that you want to create. In most cases, you will be creating a primary partition, but if you are creating more than three partitions on your disk, you will need to create an extended partition in which you can subsequently create logical partitions.
- **Filesystem:** The type of filesystem that you want to create on the selected partition. Choices are ext2, ext3, fat16, jfs, linux-swap, ntfs, reiserfs, xfs, or unformatted. The default option is ext2, although you will probably want to selected ext3 when creating a Linux filesystem to take advantage of some of the advanced features, such as journaling. To create a partition that you can access natively from a Microsoft Windows system, you will need to select either fat32 or ntfs. (fat16 filesystems can also be used by Windows, but they are tiny and antiquated—don't use them unless you have some specific reason to do so.) Selecting any entry other than unformatted causes the GNOME Partition Editor to format your new partition, creating a filesystem of the specified type in the new partition.
- **Label:** An optional text string that is used as the label for the partition. This can be handy if you use disk labels to identify partitions in your `/etc/fstab` file, as discussed in Chapter 4.

To create the new partition with the specified characteristics, click Add. The dialog shown in Figure 24-38 is redisplayed, but it has been updated to reflect the new partition that you have specified, as shown in Figure 24-41.

The pane at the bottom of Figure 24-41 lists scheduled operations for your disk, while the graphical and partition list information reflects what your disk will look like once those operations have been completed. In other words, your new partition is scheduled to be created, but it has not yet actually been created. This enables you to define multiple partitions and create them all at one. To proceed with any scheduled operations, click Apply. A confirmation dialog appears, asking if you're sure that you want to perform the specified operation(s). Click Apply to proceed. A progress dialog is displayed while the scheduled operations are being performed. Once the scheduled operations are completed, the GNOME Partition Editor dialog reappears, showing your new partition(s). (See Figure 24-42.)

FIGURE 24-41

Partitions and planned operations in the GNOME Partition Editor

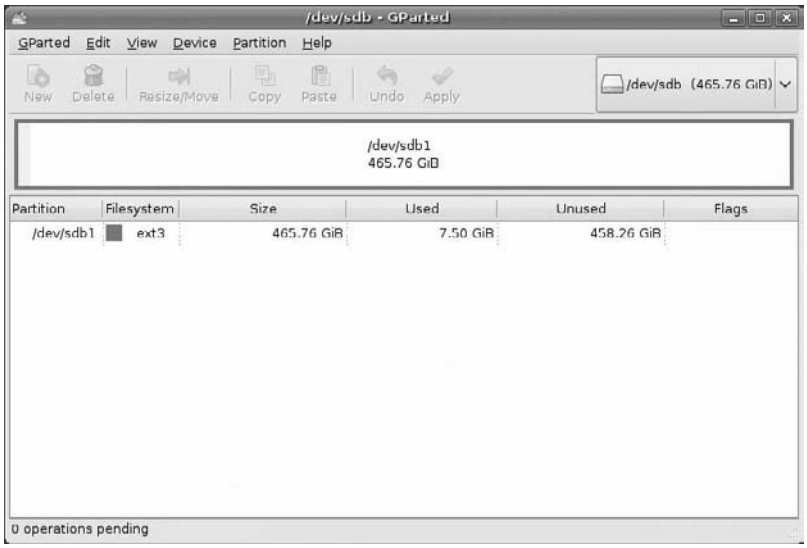
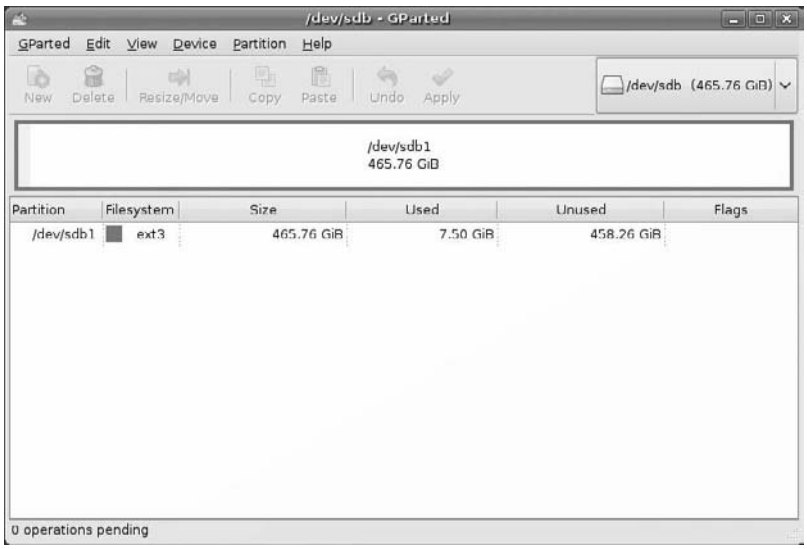


FIGURE 24-42

The GNOME Partition Editor after creating a new partition



Using Command-Line Tools to Identify Drives and Partitions

The `df` command provides an easy way to examine all of the disk partitions and associated filesystems that are currently mounted and in use on your system. Unfortunately, commands such as `df` don't provide information about devices that are newly attached to your system (and that therefore may not contain partitions yet) or devices that are attached to your system but whose filesystems are not mounted. Luckily, Linux provides some easy ways of finding all of the storage devices on your system, regardless of whether they're formatted, partitioned, mounted, or not. You'll have to put up with a small bit of cryptic command-line magic, but you're already at least an apprentice Linux wizard, right? After all, you just put a new drive in your system!

The easiest way to determine the device nodes associated with the storage devices in your system is to look through the kernel boot messages for patterns that match the strings discussed earlier in this chapter in the section entitled "Troubleshooting Automatic Device Detection." That section suggested looking through the system's log files for messages associated with hotplug events, which are the device detection and removal events that are processed by the kernel when a system is running. Luckily, you can also look through the in-memory portion of this message log by using a command called `dmesg`. Because you know that IDE devices traditionally had names that began with `hd` and are followed by some other letter, that SATA and SCSI devices have names that begin with `sd` and are followed by some letter, and that that IDE devices on newer kernels follow the same naming convention as SATA and SCSI drives, you can use a pattern such as `'[hs]d[a-z]: '` to see what messages match these device names. If you don't speak Linux regular expressions fluently, the string `'[hs]d[a-z]: '` means, "match any sequence of characters that begins with an h or an s, followed by the letter d, followed by any character between a and z (inclusive), followed by a colon, followed by a space." That was easy, wasn't it?

The result of searching for this string in the output of the `dmesg` command on a sample system is pretty much exactly what you'd hope for, as the following example shows:

```
$ dmesg | grep 'hs]d[a-z]: '
[  846.564436] SCSI device sda: 234441648 512-byte hdwr sectors (120034 MB)
[  846.564453] SCSI device sda: drive cache: write back
[  846.564514] SCSI device sda: 234441648 512-byte hdwr sectors (120034 MB)
[  846.564526] SCSI device sda: drive cache: write back
[  846.564531]  sda: sda1 sda2
[  846.575071] SCSI device sdb: 156299375 512-byte hdwr sectors (80025 MB)
[  846.575086] SCSI device sdb: drive cache: write back
[  846.578550] SCSI device sdb: 156299375 512-byte hdwr sectors (80025 MB)
[  846.578826] SCSI device sdb: drive cache: write back
[  846.578830]  sdb: sdb1 sdb2 < sdb5 >
[  848.272430] hdc: DVDRW 16X16X, ATAPI CD/DVD-ROM drive
[  848.952828] hdc: ATAPI 48X DVD-ROM DVD-R CD-R/RW drive, 2048kB Cache
[  852.169148] SCSI device sdc: 8910423 512-byte hdwr sectors (4562 MB)
[  852.170635] SCSI device sdc: drive cache: write back
[  852.171452] SCSI device sdc: 8910423 512-byte hdwr sectors (4562 MB)
[  852.172941] SCSI device sdc: drive cache: write back
[  852.172943]  sdc: sdc1 sdc2
```

On this system, these messages tell you that three storage devices using SCSI device naming conventions have been found, plus an old IDE/ATA CD/DVD-ROM drive that is available as `/dev/hdc`. The first storage device, `/dev/sda`, doesn't appear to contain any partitions. The second storage device, `/dev/sdb`, contains two physical partitions: `/dev/sdb1` and `/dev/sdb2`, and a third virtual partition, `/dev/sdb5`. (This probably means that the partition `/dev/sdb2` is an extended partition that contains the

partition /dev/sdb5—more about that in a moment.) The third storage device, /dev/sdc, contains two physical partitions: /dev/sdc1 and /dev/sdc2.

You can use the fdisk command and its -l (list) option to get more detailed information about any of these storage devices. For example, fdisk provides the following information about the partitions on it:

```
$ sudo fdisk -l /dev/sdb
Disk /dev/sdb: 80.0 GB, 80025280000 bytes
255 heads, 63 sectors/track, 9729 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot   Start    End     Blocks   Id  System
/dev/sdb1   *        1   9566   76838863+  83  Linux
/dev/sdb2         9567   9729    1309297+   5  Extended
/dev/sdb5         9567   9729    1309266   82  Linux swap / Solaris
```

This shows that, as suspected, the partition /dev/sdb2 is an extended partition, which is a physical partition that can contain other, virtual partitions, and the partition /dev/sdb5 is a partition that was created inside that extended partition.

Your mind may be a bit numb after the preceding combination of arcane command-line syntax and kernel log messages, but this section has given you some extremely powerful insights into your system and the storage devices that it contains. Because one of these storage devices, the SATA drive /dev/sda in the first example in this section, didn't contain any partitions, it is still quite useless to your system. The next section explains how to partition and format a newly added disk from the command line.

Using Command-Line Utilities to Partition and Format Drives

Chapter 4 introduced basic computer system terms such as disks, partitions, and filesystems. As discussed in this chapter, you cannot use the storage on a disk until it has been partitioned and each of those partitions has been formatted for use, typically as a filesystem. This section explains how to actually do that using the command-line commands fdisk, mkfs, and (occasionally) mkswap. These commands do the following:

- fdisk: Creates and manipulates partitions on a specified disk. We all have our own ideas about what the *f* standards for, but it's probably *format*.
- mkfs: Creates a filesystem of a specified type on a specified partition.
- mkswap: Formats a specified partition for use as a Linux swap partition, which is a special partition that your Ubuntu system can use as virtual memory.

Once you've added a new disk to your system, the first thing that you'll want to do is partition it. The previous section used the fdisk command to simply list the available partitions on a specific storage device—now let's actually make it do some work. The following example creates two partitions on the disk /dev/sdb.

To partition a disk using fdisk, do the following:

1. Use the sudo command to start the fdisk command, supplying the name of the disk that you want to partition as an argument, as in the following example:

```
$ sudo fdisk /dev/sdb
The number of cylinders for this disk is set to 14593.
```

```
There is nothing wrong with that, but this is larger than 1024,
and could in certain setups cause problems with:
1) software that runs at boot time (e.g., old versions of LILO)
2) booting and partitioning software from other OSs
   (e.g., DOS FDISK, OS/2 FDISK)
Command (m for help):
```

The last line of the sample output is `fdisk`'s command prompt, which shows that it is waiting for a command.

2. Use the `p` (print) command to display a list of any existing partitions that may already be on the disk. After all, you don't want to accidentally try to create partitions on the wrong disk or destroy any existing data. This command and its output look like the following:

```
Command (m for help): p
Disk /dev/sda: 120.0 GB, 120034123776 bytes
255 heads, 63 sectors/track, 14593 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
   Device Boot      Start         End      Blocks   Id  System
Command (m for help):
```

Good. As the total lack of partitions shows, this disk does not currently contain any partitions.

3. Use the `n` (new) command to define a new partition:

```
Command (m for help): n
Command action
   e   extended
   p   primary partition (1-4)
```

4. Entering this command prompts you for the type of partition that you want to create, either `e` (extended) or `p` (physical). Use the `p` command to create a physical partition:

```
p
Partition number (1-4):
```

5. Specifying that you want to create a physical partition prompts you for the number of the partitions that you want to create. Linux disks can contain only four physical partitions, one of which can be an extended partition. Create partition number 1 by entering `1`, as in the following example:

```
1
First cylinder (1-14593, default 1):
```

6. Identifying the number of the physical partition that you want to create prompts you for the physical disk block at which the partition begins, and displays the range of possible values (1–14593, in this case). Press Return to accept the default value of `1`:

```
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-14593, default 14593):
```

7. When creating partitions, you can define them in terms of the disk blocks at which they begin or end, or you can define them in terms of the blocks at which they begin and the size of the partition. Enter `+30G` to indicate that you want to create a 30GB partition:

```
Last cylinder or +size or +sizeM or +sizeK (1-14593, default 14593): +30G
Command (m for help):
```

8. This completes the information required to create the partition, but doesn't tell you much about whether this actually succeeded. Use the `p` command again to list any partitions that are currently defined:

```
Command (m for help): p
Disk /dev/sda: 120.0 GB, 120034123776 bytes
255 heads, 63 sectors/track, 14593 cylinders

Units = cylinders of 16065 * 512 = 8225280 bytes
   Device Boot      Start         End      Blocks   Id  System
/dev/sda1               1        3648    29302528+  83  Linux
```

Sure enough, a 30GB partition has been defined as the partition `/dev/sda1`, the first partition on your disk.

TIP **At this point, no physical changes to the partitioning of your disk have been made. Any proposed changes are only in memory and will be discarded unless you write your changes to disk before exiting the `fdisk` command.**

9. Repeat this process to create a second, smaller partition: No play-by-play this time—the complete sequence of prompts and commands used to do this is the following:

```
Command (m for help): n
Command action
   e   extended
   p   primary partition (1-4)
p
Partition number (1-4): 2
First cylinder (3649-14593, default 3649):
Using default value 3649
Last cylinder or +size or +sizeM or +sizeK (3649-14593, default 14593): +1G
```

10. This completes the information required to create the second partition, but again, doesn't tell you much about whether it actually succeeded. Use the `p` command again to list any partitions that are currently defined:

```
Command (m for help): p

Disk /dev/sda: 120.0 GB, 120034123776 bytes
255 heads, 63 sectors/track, 14593 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
```

```
   Device Boot      Start         End      Blocks   Id  System
/dev/sda1               1        3648    29302528+  83  Linux
/dev/sda2            3649        3771      987997+  83  Linux
 Command (m for help):
```

Sure enough, you now have both a 30GB partition and a second 1GB partition.

11. Now, write these changes to disk and exit the `fdisk` command, because these examples are sufficient for my purposes in this section. You use the `w` (write) command to write partition changes to the disk and exit the `fdisk` command:

```
Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
```

Congratulations, you've just manually partitioned a disk. Normally, I wouldn't have left unused space on this disk, but I'm only doing this as an example, and I can always rerun `fdisk` to create additional partitions in the remaining space.

The next step is to format each defined partition either as a filesystem or as additional swap space. Graphical utilities usually do this for you when creating partitions, but things aren't quite as user-friendly from the command line. To be honest, you'll very rarely create additional swap partitions after you've initially installed your system, but I want to make sure that you know how to do it just in case. 99.9% of the time, you'll be creating filesystems on the partitions of the disks that you add to your system, so that you can mount them to increase the amount of available storage on your system.

Creating a filesystem is done using the `mkfs` utility, which is actually just a wrapper for other commands that create different types of filesystems. On most Linux systems, you will probably be creating filesystems in any of the following formats: ext2, ext3, JFS, ReiserFS, Windows Virtual FAT (vfat), or XFS. These types of filesystems were explained in the section of Chapter 4 entitled "Local Filesystems: Standard and Journaling." To create any of these types of filesystem using the default options for that type of filesystem, you simply specify the type of filesystem that you want to create using the `mkfs` command's `-t` (type) option. The following example shows the command and output used to create an `ext3` filesystem on the partition `/dev/sda1`:

```
$ sudo mkfs -t ext3 /dev/sda1
mke2fs 1.41.3 (12-Oct-2008)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
3662848 inodes, 7325632 blocks
366281 blocks (5.00%) reserved for the super user
First data block=0
224 block groups
32768 blocks per group, 32768 fragments per group
16352 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632,
        2654208, 4096000
```

```
Writing inode tables:   224/224
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 25 mounts or
180 days, whichever comes first. Use tune2fs -c or -i to override.
```

That's all there is to it! You can now mount this partition temporarily by using the `mount` command, or permanently add it to the list of filesystems that will be mounted each time you boot your system. Both of these concepts were explained in Chapter 4.

The `mkswap` command enables you to format a specified partition in a special way so that it can be used as swap space for use by your system's virtual memory manager. As an example, the command used to create a swap partition on `/dev/sda2` and the output from that command are the following:

```
# mkswap /dev/sda2
Setting up swapspace version 1, size = 1011703 kB
no label, UUID=731c0f53-bf25-4f91-9ac5-73a2b5bef1eb
```

You can begin using this swap space immediately by executing the command `swapon /dev/sda2`. To add this swap space to your system permanently, add it to your system's `/etc/fstab` file, as described in Chapter 4.

Using PCMCIA Cards

PCMCIA (Personal Computer Memory Card International Association) cards are the original mechanism designed to support removable and insertable interfaces in portable computer systems. PCMCIA cards are now often just called PC Cards because no one can ever remember what the acronym means (as suggested by the optional PCMCIA expansion, "People Can't Memorize Computer Industry Acronyms"). Thanks to the rise of USB, PCMCIA cards aren't as common as they once were, but they're still a standard feature in many laptops.

PCMCIA cards are available that provide wired and wireless network interfaces, modems, SCSI, USB, and FireWire interfaces for computer systems that don't have them, and so on. With the 2.6 kernel used on Ubuntu systems, the traditional Linux PCMCIA support package has been integrated into the hotplug system to support the creation of appropriate device nodes when PCMCIA cards are inserted.

The primary application to identify any PCMCIA cards inserted in your system is the `pccardctl` application (formerly known as `cardctl`). However, changes to how the kernel handles PCMCIA devices mean that this utility may not return detailed information about a PCMCIA card. To see what cards are present in your system and to display summary information about them, execute the `pccardctl ident` command. That will return something like the following:

```
$ pccardctl ident
Socket 0:
  no product info available
Socket 1:
  product info: " LEXAR ATA FLASH CARD    ", " STORM ", "TE01"
  manfid: 0x4e01, 0x0200
  function: 4 (fixed disk)
```

This shows that I have two PCMCIA cards inserted but that no information is available about the first one. Some of the other `pccardctl` commands may provide more information. The `pccardctl ls` command, for instance, displays the following:

```
$ sudo pccardctl ls
Socket 0 Bridge:      [yenta_cardbus]    (bus ID: 0000:02:03.0)
  CardBus card -- see "lspci" for more information
Socket 1 Bridge:      [yenta_cardbus]    (bus ID: 0000:02:03.1)
Socket 1 Device 0: [pata_pcmcia]   (bus ID:  1.0)
```

This shows that the PCMCIA network card in my first PCMCIA slot is being handled as a PCI card, while the compact flash card is being handled by the Parallel IDE (pata) subsystem.

In most cases, the hotplug/udev subsystems will automatically create necessary device nodes, using the standard mechanisms to notify you when new PCMCIA storage devices are attached, when new PCMCIA wireless interfaces appear, and so on.

Adding PCI Cards

PCI cards are slightly different animals than most of the devices I've discussed elsewhere in this chapter, because they typically provide support for other devices, such as speakers (PCI sound cards), and storage devices (PCI USB, FireWire, or SCSI cards), or support for basic devices such as modems, additional serial or parallel ports, and so on. Your system must be down in order to add a PCI card, and discovering one during the boot process isn't a notable event. The information that may appear in the kernel log related to new PCI cards is the discovery of devices that are attached to them, such as a modem, serial ports, disk drives, and so on.

Nevertheless, it is useful to be able to list the PCI devices found in your system, to verify that cards have been correctly inserted and identified by the system. The command used to do this is the `lspci` command, which displays output like the following:

```
$ lspci

0000:00:00.0 RAM memory: nVidia Corporation C51 Host Bridge (rev a2)
0000:00:00.1 RAM memory: nVidia Corporation C51 Memory Controller 0 (rev a2)
0000:00:00.2 RAM memory: nVidia Corporation C51 Memory Controller 1 (rev a2)
0000:00:00.3 RAM memory: nVidia Corporation C51 Memory Controller 5 (rev a2)
0000:00:00.4 RAM memory: nVidia Corporation C51 Memory Controller 4 (rev a2)
0000:00:00.5 RAM memory: nVidia Corporation C51 Host Bridge (rev a2)
0000:00:00.6 RAM memory: nVidia Corporation C51 Memory Controller 3 (rev a2)
0000:00:00.7 RAM memory: nVidia Corporation C51 Memory Controller 2 (rev a2)
0000:00:02.0 PCI bridge: nVidia Corporation C51 PCI Express Bridge (rev a1)
0000:00:03.0 PCI bridge: nVidia Corporation C51 PCI Express Bridge (rev a1)
0000:00:04.0 PCI bridge: nVidia Corporation C51 PCI Express Bridge (rev a1)
0000:00:05.0 VGA controller: nVidia Corporation C51 PCI Express Bridge
0000:00:09.0 RAM memory: nVidia Corporation MCP51 Host Bridge (rev a2)
0000:00:0a.0 ISA bridge: nVidia Corporation MCP51 LPC Bridge (rev a2)
0000:00:0a.1 SMBus: nVidia Corporation MCP51 SMBus (rev a2)
0000:00:0b.0 USB Controller: nVidia Corporation MCP51 USB Controller (rev a2)
0000:00:0b.1 USB Controller: nVidia Corporation MCP51 USB Controller (rev a2)
0000:00:0d.0 IDE interface: nVidia Corporation MCP51 IDE (rev a1)
```

```
0000:00:0e.0 IDE interface: nVidia Corporation MCP51 Serial ATA Controller
0000:00:0f.0 IDE interface: nVidia Corporation MCP51 Serial ATA Controller
0000:00:10.0 PCI bridge: nVidia Corporation MCP51 PCI Bridge (rev a2)
0000:00:10.1 0403: nVidia Corporation MCP51 High Definition Audio (rev a2)
0000:00:14.0 Bridge: nVidia Corporation MCP51 Ethernet Controller (rev a1)
0000:00:18.0 Host bridge: HyperTransport Technology Configuration
0000:00:18.1 Host bridge: [Athlon64/Opteron] Address Map
0000:00:18.2 Host bridge: [Athlon64/Opteron] DRAM Controller
0000:00:18.3 Host bridge: K8 [Athlon64/Opteron] Miscellaneous Control
0000:04:05.0 FireWire (IEEE 1394): VIA, Inc. IEEE 1394 Host Controller
0000:04:08.0 SCSI storage controller: LSI Logic / Symbios Logic 53c895
0000:04:09.0 Ethernet controller: 3Com Corporation 3c595 100BaseTX [Vortex]
```

Although this output is somewhat ugly, it is easy enough to scan to determine if your new card has been detected. If it has, you should proceed to the configuration utility associated with the capability or type of device that the PCI card provides.

Examining and Troubleshooting Devices Graphically

Ubuntu and Kubuntu systems provide a great graphical utility called the Device Manager that enables you to browse, navigate through, and explore the hardware on your system from your system's point of view. The Ubuntu and Kubuntu versions of this utility aren't installed by default, but you can install them using your favorite package management software (as discussed in Chapter 23) by installing the `gnome-device-manager` package on a Ubuntu system or the `kde-hal-device-manager` package on a Kubuntu system. Once it is installed, select Applications ➪ System Tools ➪ Device Manager on a Ubuntu system to start this tool, which displays the initial dialog shown in Figure 24-43.

After installing the `kde-hal-device-manager` package on a Kubuntu system, start it by selecting the K Menu ➪ System ➪ Device Manager ➪ Hardware Viewing menu command on a KDE3-based Kubuntu 8.04 or earlier system, or the K Menu ➪ Applications ➪ System ➪ Device Manager ➪ Hardware Viewing menu command on a KDE 4-based Kubuntu 8.10 or later system. Figure 24-44 shows the Device Manager on a Kubuntu 8.10 system; it looks essentially the same on a KDE 3-based Kubuntu system.

Once the Device Manager is displayed, you can use the pane at the left to scroll or drill down to examine any of the basic devices provided by your system or attached to any of its interfaces. As you select different devices, appropriate tabs display in the Device Manager's right pane. The right pane typically provides a general Device information tab and an Advanced tab, which provides additional details about a selected device or interface.

The information that this utility provides can be quite useful in determining how—or if—your system's hardware is being recognized, which can help you in terms of knowing where to start when trying to resolve problems with new hardware. Seeing how your system recognizes, organizes, and uses hardware is also just plain interesting to many of us.

FIGURE 24-43

Ubuntu's Device Manager

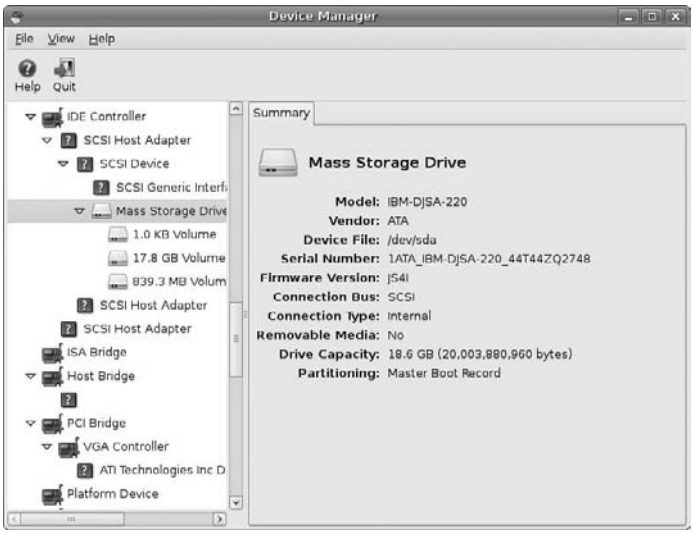
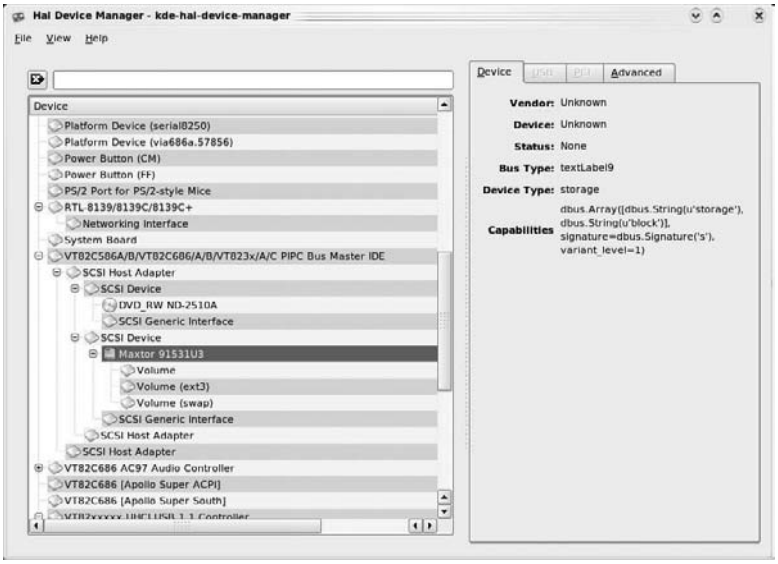


FIGURE 24-44

The Device Manager on a Kubuntu system



Summary

This chapter explained how to add various types of hardware to your system, ranging from removable storage that you want to use only temporarily, to more permanently attached devices like printers, new internal hard drives, and scanners.

Chapter 25 explores how to secure your Ubuntu system. Though Linux is generally immune to most of the types of viruses, Trojans, and spyware that infest the Internet (and most Microsoft Windows systems), network security is as critical on a Linux system as it is on any other type of computer system that may be directly connected to the Internet. Hundreds of books are available that are dedicated to network and computer security, but Chapter 25 provides a basic explanation of the key concepts and concerns, explaining basic ways in which you may want to secure your system and highlighting areas that you may want to pursue further.

Chapter 25

Network Configuration and Security

Almost from the very beginning of home computing in the 1970s, personal computers have reached out to touch other types of computer systems. Long before ISPs, and before the Internet even existed, home computer fans used modems to access bulletin board systems, remote mainframe or minicomputers, and ancient content providers like Compuserve and AOL, using various terminal emulation programs to communicate with each other, transfer files, and so on. Early store-and-forward mechanisms such as the Unix-to-Unix Calling Program (UUCP) and fidonet provided great ways of disseminating files and other information across slow networks of computer systems that were networks only in the sense that they knew each other's phone numbers.

The conversion of the ARPANET to the Internet and its resultant commercialization gave birth to the notion of ISPs, commercial Internet service providers, who provided a mechanism for home computers to directly access the Internet, albeit through kludgey point-to-point solutions that still depended on a modem and thus provided Net surfing speeds that were only guaranteed (supposedly) to be greater than zero. Regardless, the advent of the ISP ended the concept of the PC as an asynchronous island, making it a real participant in the Internet, even if slowly.

As ISPs surfaced and became a fundamental utility for many home computer users, networking and PC hardware costs continued to drop, approaching the commodity hardware pricing normally associated with toasters and refrigerators. The reality of more and more home computer users, even in the same homes, introduced the notion of home computer networks, often standalone or with modems still connecting specific systems to the Internet by functioning as a 9600 or 56KB gateway to the Internet via an ISP. We all owe much to those pioneering users of home computer who were willing to access the Net and download porn though such tragically slow connections.

Broadband Ethernet, even cheaper wired network hardware, and the explosive growth of wireless networking has made networking a true reality for many home computer users. Home computer systems may now have real IP addresses

IN THIS CHAPTER

Graphical tools for network configuration

Traditional network configuration files

Testing network interfaces

Tips for securing your network

Installing a firewall

and functional connection speeds to the Internet, and are also commonly members of home computer networks that share those connections to the Net using mechanisms such as Native Address Translation (NAT). If you're new to networking, the first section of this chapter provides an overview of basic networking concepts and terminology.

Better networking and network access comes at a price. Ubiquitous networking gives thousands of "randoms" access to your computer system through a real IP address or Web server and other network processes. Most of them could care less, some are simply curious, and others are downright malicious. The last set gives everyone else a bad name by actively trying to break into computer systems to exploit them in some fashion. I have no problem with hackers who are simply curious about what's out there—exploration has always been a fundamental part of the human condition. Real *hackers* are the electronic equivalent of the National Geographic Society or *Star Trek*, boldly going where no person has gone before. People who break into systems to damage or exploit them are *crackers* (my apologies to readers in the southern parts of the United States), who give everyone else a bad name.

Unfortunately, there are plenty of unscrupulous crackers who would love to break into your machine and damage it or turn it into some sort of zombie system, either to supposedly demonstrate their cleverness or to somehow make a buck. Sigh. Ubiquitous networking begets the easy availability of tools that enable this sort of thing. The people that use them are often so-called "script kiddies," who use existing tools to demonstrate cleverness in the same way that giving a child an Uzi demonstrates marksmanship.

The bottom line of ubiquitous networking is that security becomes everyone's job. If you live in a small town that considers taking two newspapers from the box on the corner a serious crime, locking your door at night may seem silly. Unfortunately, when you use a personal computer with network access, you are part of the big city known as the Internet. The administrators of enterprise and academic systems that require continuous access to the Internet have known this for a long time. Sadly enough, nowadays your grandmother, parents, and you have to worry about it, too. Security is more of a concern today than it has ever been before, and tomorrow will just be worse.

This chapter provides a basic introduction to networking, explains the tools that Ubuntu and Kubuntu Linux provide to manually and graphically configure and test your network, and (most importantly) provides some general guidelines on how to secure your system to protect it as best as anyone can. There's an old saying in the IT biz that the only truly secure system is one that isn't connected to anything. Although this is true, it's also impractical. There are easy rules to follow to minimize the chances that your system will be broken into. You're already running Ubuntu or Kubuntu, which puts you miles ahead of the millions of vulnerable Microsoft Windows users out there.

Networking 101

Most modern computer systems can communicate with other systems and devices over a type of network called Ethernet, using the Transmission Control Protocol/Internet Protocol (TCP/IP) and Universal Data Packet (UDP) protocols. Ethernet was invented by Xerox Corporation at Xerox PARC (Palo Alto Research Center) in the early 1970s. Like most things they've invented—except for the photocopier—Xerox failed to make money from Ethernet, which was actually commercialized by many companies (like 3COM, which was founded by the inventor of Ethernet networking, Bob Metcalf, who knew a good thing when he invented it).

Until a few decades ago, "the Internet" was a fairly techie term, used only by people whose employers or academic experience offered connectivity to the Internet or its predecessor, the ARPANET. The

creation and popular explosion of the World Wide Web and the advent of e-mail as a replacement for phone calls changed all that—suddenly, there was a reason for people to want (or perhaps even need) access to the Internet.

Early home Internet connectivity was primarily done through dial-up connections that emulated TCP/IP connections over dial-up lines using protocols such as Serial Line Internet Protocol (SLIP), Compressed SLIP (CSLIP), or Point-to-Point Protocol (PPP). Unless you were a serious computer geek, developer, or researcher, a home network was somewhat rare, but the advent of broadband access to the Internet through cable and telephone providers changed all that. As mentioned in the introduction to this chapter, home networks are becoming more common, but most people have never needed to set one up before now. If you use a single PC, Mac, or workstation as your sole home machine, a straight connection to a cable or DSL modem works just fine. However, the instant you want to enable multiple machines to communicate over a home network, you may encounter unfamiliar terms like hubs, switches, 10-BaseT, RJ-45, crossover cables, uplink ports, packets, gateways, routers, Cat 5, and a variety of others that pass for popular nouns among nerdier users. This section provides a quick overview of these terms. It tells you how to set up a simple home network and makes you comfortable with the network-related terms that are used throughout this book. For more detailed information, consult any of the hundreds of books available on home networking.

The basic element of a modern network connection is a standard Ethernet cable, which is just a length of multistrand cable with connectors on either end that enable you to connect a network card in your personal computer (or whatever type) to another network device. The most common connectors used today are plastic connectors known as RJ-45 connectors, which is a transparent plastic jack that looks like a fatter version of a standard telephone cable connector. Ethernet cables that use these connectors are often known as 10-BaseT, 100-BaseT, or even 1000-BaseT, where the numeric portion of the name indicates the speed of your network—the cables are the same. 1000-BaseT is more commonly known as Gigabit Ethernet, and is the up-and-coming standard, because things tend to get faster. 10/100 Ethernet (10 megabit or 100 megabit) is the standard nowadays.

NOTE **You may also encounter the term 10-Base2 when researching network cards. This is an older type of 10-megabit Ethernet cabling that uses shielded Bayonet Neill-Concelman, or Baby N Connector (BNC) cables, and is not supported by most networking hardware today.**

The best way to visualize the Internet or any Ethernet network is as an extremely long piece of cable to which several computers and network devices are attached. In the simplest case, you must use a device called a hub, switch, or router to attach multiple machines to an Ethernet. A *hub* is a device with multiple incoming connectors for attaching the Ethernet cables from different machines, with a single output connector that attaches it to another Ethernet device such as a cable modem, another hub, or a switch, router, or gateway. Network communications on any incoming port of the hub are broadcast to all other devices on the hub and are also forwarded through the outgoing connection. Switches are much like hubs on steroids because they keep track of how network connections between different machines are made and reserve dedicated internal circuitry for established connections. Switches are, therefore, both typically faster and more expensive than hubs because they do more.

Gateways and routers are similar to hubs and switches, but are designed to provide connectivity between different networks. If a machine that you are trying to connect to isn't immediately found on your local network, the request is forwarded through your gateway, which then sends it on. Network communication is done using discrete units of information that are known as *packets*. Packets contain the Internet Protocol (IP) address of the host that they are trying to contact. IP addresses are in the form of *NNN .NNN.NNN.NNN* and are the network equivalent of a post office box, uniquely identifying a specific machine. Packets for an unknown local host are sent through your gateway. Routers are expensive,

sophisticated pieces of hardware that direct network communication between multiple networks, translate packets between different network communication protocols, and limit network traffic to relevant networks so that your request to retrieve a file from your local music server isn't broadcast to every machine on the Internet.

IP addresses use additional network configuration data, known as a *netmask*, to identify the portion of a network to which a given machine belongs. The netmask identifies the bits in your IP address that identify a network segment, and is typically 255.255.255.0 for a standard Class C IP address on a home network. The netmask 255.255.255.0 is 11111111.11111111.11111111.00000000 in binary, which means that the last 24 bits in a binary IP address (i.e., the numbers after the final period in a standard IP address) are used to identify a specific machine. Some ISPs use more complex netmasks to divide their networks into subnets that provide better control over the IP addresses that they issue and the associated network traffic.

The most common way to connect machines on a home network is to use a hub or a home gateway that is connected to your cable or DSL modem. The difference between these is that a hub simply forwards packets through its outgoing connector (known as an *uplink port* because it simply links the network connections on that device with those on another, forwarding network packets to the other device and is, therefore, wired differently). A home gateway may convert internal network addresses to addresses that are compatible with the outside world before sending the information on through its outgoing or uplink connector. If you're using a hub to connect your home network to your cable or DSL modem, each machine on your home network would require an IP address that is unique on the Internet. This can be expensive, because most ISPs charge money for each unique host that can be connected to the Internet from your home at any given time. Home gateways provide a way around this because they enable your home network to use a special type of IP address, known as a nonroutable IP address, to assign unique internal network addresses. As the name suggests, nonroutable IP addresses are addresses that will not be forwarded (routed) by the networking devices that support the actual Internet. When using nonroutable IP addresses, whatever gateway you (or your ISP) use to connect to the Internet internally translates these to valid external addresses whenever you try to connect to a machine on the Internet. The most common nonroutable IP addresses are in the form of 192.168.*X*.*Y*, where *X* and *Y* are specific to how you've set up your network.

TIP **If you're really interested, you can get more information about nonroutable IP addresses and address translation in the Internet RFCs (Requests for Comment) that defined them, 1597 and 1918. Use your favorite Internet search engine to find relevant information, or check out links such as** `www.safety.net/sum1597.html` **and** `www.howstuffworks.com/nat2.htm`**.**

IP addresses are assigned to computer systems in two basic ways, either statically or dynamically. Static addresses are unique to your home network and are always assigned to a particular machine. Dynamic addresses are addresses that are automatically assigned to a computer system or network device when you turn it on. Most ISPs use dynamic addresses because only a limited number of IP addresses are available on the Internet. Using dynamic IP addresses enables your ISP to recycle and reassign IP addresses as people turn their machines off and on. Most dynamic IP addresses nowadays are assigned using a protocol called *Dynamic Host Configuration Protocol (DHCP)*, which fills out the network information for your system when it activates its network interface, including things like the IP address of a gateway system and the IP addresses of Distributed Name Service (DNS) servers that translate between hostnames and the IP addresses that they correspond to.

To use static addresses on your home network, you simply assign each network interface on each machine a unique, nonroutable IP address from a given family of nonroutable IP addresses. For example, most of my home machines have static addresses in the form of 192.168.6.*Y*. Because I use a home gateway, I've

configured it to do address translation (more specifically known as *NAT*, or *Network Address Translation*) to correctly translate between these addresses and the external IP address of my home gateway box.

TIP **Systems on which multiple Ethernet connections are available are quite common today. If you are using multiple Ethernet connections simultaneously, it usually only makes sense to have each interface connected to a different network, because network routing is somewhat confusing otherwise. Systems with multiple Ethernet connections where each of these connections is attached to different networks are known as *multi-homed* systems, because they have an identity (home) on each network.**

If you want to use Dynamic IP addresses on a home network, one of the machines on your home network must be running a DHCP server. Most home gateways, such as those from DLink or Linksys, have built-in DHCP servers that you simply configure to hand out IP addresses from a specific range of addresses (192.168.6.200 through 192.168.6.250, in my case). Once you activate address translation on your home gateway, your gateway will route packets appropriately. Remember that your home gateway is probably getting its IP address by contacting your ISP's DHCP server, whereas hosts on your internal network will get their IP addresses from your DHCP server. Don't set up hosts on an internal network to contact your ISP's DHCP server unless you have only a single machine on your home network or want every one of your machines to be visible on the Internet. If you are using a home gateway that doesn't provide a DHCP server, want to have more control over what your DHCP server does, or are using Ubuntu Linux in an enterprise or commercial setting, you may want to set up your own DHCP server on an Ubuntu system, which is explained in Bonus Chapter 10, "Setting Up a DHCP Server."

A final aspect of networking is how your system identifies and locates specific computer systems on the Internet. This is typically done through the Domain Name Service (DNS), which is explained in the introduction to Bonus Chapter 11 and the section of that chapter entitled "Overview of DNS and BIND."

The overview in this section should have familiarized you with basic general and home networking terms and concepts. As you might expect, the Internet is knee-deep in Web sites that provide more general information about home networking. For truly detailed information about setting up and configuring a home network on a specific type of machine and operating system, see any of the hundreds of books on those topics at your local bookstore.

Configuring Your Network Hardware

Configuring the network hardware on your computer system is part of the Ubuntu or Kubuntu installation process, both of which use an automatic network configuration mechanism known as the Network Manager (`http://projects.gnome.org/NetworkManager`). However, things change. You may install new network hardware, change existing hardware from relying on DHCP to using static IP addresses on your network, or simply want to have a better understanding of how networking works or is configured on your system(s). The Network Manager is designed to auto-detect networking hardware and network configuration changes, which can be especially handy when you're using a laptop and want to be able to connect your system to the Internet from your home, from random Wi-Fi hotspots, and perhaps even from work. The Network Manager interacts directly with your hardware using HAL, the Linux Hardware Availability Layer, and D-Bus, a message bus system that simplifies communication between applications regarding resources, system events, and so on. These provide a great deal of flexibility but can make it difficult to debug configuration problems, especially if you're used to the traditional network configuration mechanisms used on Debian-based distributions such as Ubuntu.

The next two sections discuss how to use the standard Network Manager package that's installed by default on Ubuntu and Kubuntu systems. As always, there are some differences between how you interact with the Network Manager on Ubuntu and Kubuntu systems, and then again between KDE 3-based Kubuntu 8.04 and KDE 4-based Kubuntu 8.10 or later systems.

No matter how cool the Network Manager is, it doesn't work for everybody. Not to worry! The sections on the Network Manager are followed by a section that discusses how to configure your networking hardware using the traditional Debian/Ubuntu mechanism, which simply involves editing a text file. This may be best approach for you if you have problems with the Network Manager or simply don't care about graphical tools for network configuration. As always, Linux is knee-deep in options and alternatives, and the choice is yours.

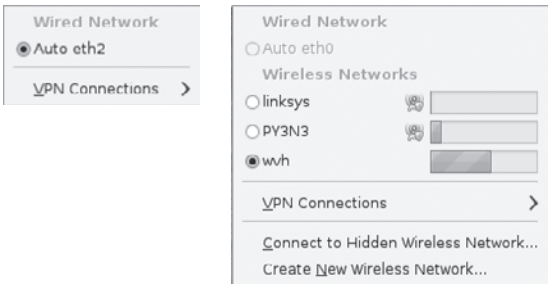
TIP **This section provides information about using the graphical tools provided with Ubuntu and Kubuntu systems for configuring both wired and wireless connections over Ethernet devices that use standard Linux device drivers. If you're having problems configuring your wireless connections using the Network Manager, see Chapter 26, "Going Wireless," for more detailed information about wireless networking and a discussion of other tools that you may find useful when configuring wireless networks. This includes command-line tools and a package called `Wicd`, an open source wired and wireless network manager that is an excellent (some say superior) alternative to the standard Network manager. Chapter 26 also explains how to use Microsoft Windows drivers on your Ubuntu or Kubuntu system if Linux does not yet support your Ethernet hardware. This is, unfortunately, somewhat common with wireless networking hardware—not all Ethernet hardware manufacturers have the same level of enlightenment.**

Using Network Manager on Ubuntu

By default, all interaction with network devices in Ubuntu 8.10 is done through the Network Manager's GNOME panel applet, `nm-applet`. This applet is displayed in the Notification Area in the top right portion of the GNOME panel, using an icon that looks like a frontal view of two slightly offset computer screens if you are connected to a wired network or a histogram if you are connected to a wireless network. (A red "x" is displayed at the bottom of the two-screens icon if you are not currently connected to a network.) Left-clicking this icon displays a menu that lists the current status of all of the network connections that were detected on your system, as shown in Figure 25-1.

FIGURE 25-1

Drop-Down status information from the Network Manager applet

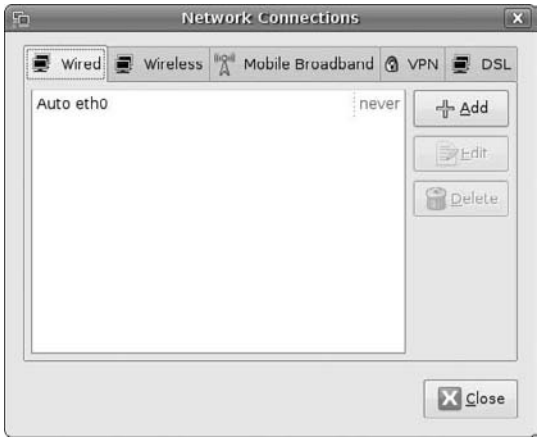


The left portion of Figure 25-1 shows the Network Manager applet's drop-down menu on a system that has a single wired Ethernet interface that is recognized as the Ethernet device `eth2`. The selected radio button to the left of the name of that interface shows that it is currently active. The right portion of Figure 25-1 shows the Network Manager applet's drop-down menu on a system that has a single wired Ethernet interface that is recognized as the Ethernet device `eth0`, and that also has a wireless Ethernet interface that has detected the availability of three wireless networks. The wired Ethernet interface is grayed out, indicating that it is not in use, and the selected radio button to the right of the wireless network `wvh` shows that the system's wireless Ethernet interface is currently connected to that wireless network. The bars to the right of each wireless network indicate the strength of the signal from an access point for that wireless network. The icon to the left of two of those bars indicates that these are secure networks that require authentication in order to access them, which the network named `wvh` does not.

To configure your network connections, you can either select the System ➪ Preferences ➪ Network Configuration menu item or right-click on the Network Manager applet and select Edit Connections to display the dialog shown in Figure 25-2.

FIGURE 25-2

The Network Manager applet's Network Connections dialog



The Network Connections dialog displays your wired connections by default. Network connections that your system will try to automatically establish each time you restart the system are preceded by the word "Auto".

The next three sections explain how to configure wired, wireless, and VPN connections using the Network manager applet on your Ubuntu system. Remember that the Network Manager is designed to automate enabling and configuring the network interfaces on your system. In most cases, you won't have to actually do anything for standard wired and wireless Ethernet connections—if you turn on your Ubuntu system and its network interface(s) come up automatically and correctly, you don't have to do any additional configuration and can just start reading mail, surfing the Web, downloading porn, or any other activity that requires a network connection. However, if your interfaces didn't start automatically, use DHCP where you want to use a static IP address, you want to connect to a different wireless network, or if you want to create a connection to a virtual private network, read on.

Configuring Wired Connections on Ubuntu

To examine or modify how a specific wired connection is configured, make sure that the Wired tab is displayed, select the name of the connection that you want to configure, and click Edit. This displays a dialog like the one shown in Figure 25-3.

FIGURE 25-3

Information about a specific wired network connection



The Connect automatically option on this dialog is selected, showing that your system automatically tries to enable this connection each time it boots, and the System setting option identifies this network connection as a basic system setting rather than a user-configured setting—this network connection will be activated and available for all users of your system.

To see additional details for this network connection, select the IPv4 Settings tab to display the pane shown in figure 25-4.

The selection in the Method drop-down list on this dialog identifies how your system obtains its IP address. The most common of these are DHCP, the default, and Manual, which enables you to supply static IP address information. Static IP addresses are typically used for server systems, so that they can always be found at the same IP address—after all, you don't want the systems on your home network to have to search for your print server each time you print a file. It's much easier for print and other servers to simply use a static IP address.

FIGURE 25-4

IP address configuration for a wired network connection



To manually specify an IP address, click the Method drop-down list and select Manual. This enables the other configuration options on this dialog, as shown in Figure 25-5.

Click Add to specify the IP address for this network connection, and enter the IP address, netmask, and gateway for this connection. Next, enter the address for one or more DNS servers, separating each IP addresses from the next with a comma, and enter the name of the domain of which this machine is a member.

NOTE **Because DNS servers are the usual source of information that map IP addresses to hostnames, you should enter only IP addresses in this dialog. If you specify a hostname, your system will need to use a DNS server to figure out the IP address associated with that DNS server, which would cause a nasty chicken-and-egg loop.**

Configuring Wireless Connections on Ubuntu

To examine or modify how a specific wireless connection is configured, or to create a new wireless connection, select the Wireless tab, which displays a pane like the one shown in Figure 25-6.

TIP **If you are connecting to a wireless network that uses WEP 40/128-bit Key, WEP 128-bit Passphrase, and no security (which are almost functionally equivalent nowadays), you can quickly define a new wireless network connection by right-clicking on the Network Manager applet and selecting the Create New Wireless Network menu option. This displays a dialog that prompts you for the name of the wireless network connection that you want to create, enables you to select between WEP 40/128-bit Key, WEP 128-bit Passphrase, and no security, and enables you to enter the key or passphrase for these two common WEP flavors.**

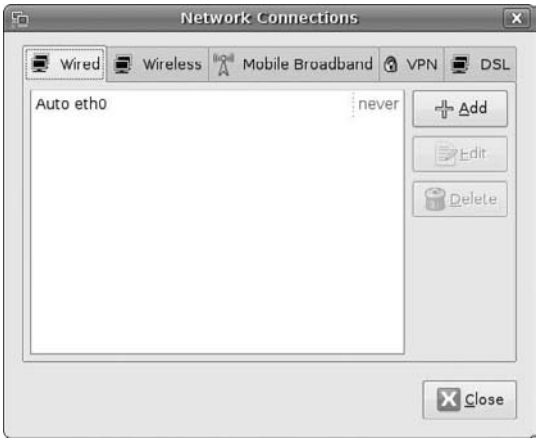
FIGURE 25-5

Manually configuring a wired network connection



FIGURE 25-6

General Information about wireless network connections



In this example, the Wireless pane shows two wireless network connections. To examine a specific connection, select the name of that connection and click Edit. To create a new wireless connection, click Add. Either of these actions displays a dialog like the one shown in Figure 25-7.

FIGURE 25-7

Information about a specific wireless network connection



To make sure that your system automatically tries to enable this connection each time it boots, make sure that the Connect automatically option on this dialog is selected. You should not enable this option for wireless connections that you will only use in specific circumstances, such as when connecting to a wireless network at your office. Similarly, if the network connection that you are editing or defining is one that should be activated and available for all users of your system, make sure that the System setting option is selected to identify this network connection as a basic system setting rather than a user-specific setting. If this is a user-specific setting, make sure that the System setting option is not selected.

Next, enter or verify the SSID (Service Set Identifier) for the wireless network you are connecting to. Make sure that the mode is set to Infrastructure unless you are establishing a direct connection between the wireless interfaces on two machines. Finally, setting the BSSID (Basic Service Set Identifier) or Mac address is only necessary if you want to make sure that your system only connects to a specific wireless access point and the wireless network with which it is associated.

To proceed with examining, modifying or creating the definition of a wireless connection, click the Wireless Security tab. You can click the drop-down Security list to see your choices for the mechanism

that you system will use to authenticate itself to the wireless network that is identified on the Wireless tab. If you are connecting to a wireless network that does not use security, select None. Otherwise, select one of the following options, each of which requires security information such as a key, passphrase, certificates, or username and password combination that you will need to obtain from the administrator of the wireless network that you want to connect to:

- **WEP 40/128-bit Key:** Enables you to enter an encrypted key that will be used to authenticate the computer to the wireless network. Though quite common, WEP (Wired Equivalent Privacy) keys can be cracked in a few minutes using popular open source packages such as Aircrack-ng (`www.aircrack-ng.org/doku.php`) and Kismet (`www.kismetwireless.net/index.shtml`). Nowadays, WEP is primarily used to indicate that you are interested in security as opposed to actually providing any, and is often the equivalent of painting a big target on your network.
- **WEP 128-bit Passphrase:** Enables you to enter a passphrase that will be used to generate WEP keys that will enable you to authenticate to the wireless network. 128-bit WEP security is much better than 40-bit WEP, requiring up to an hour to crack with the same open source tools. Ask the nearest fifth grader for more details.
- **LEAP:** Cisco's Lightweight Extensible Authentication Protocol (LEAP) supports username and password-based authentication between a wireless client and a RADIUS (Remote Authentication Dial-In User Service) server such as Cisco ACS or FreeRADIUS (`http://freeradius.org`).
- **Dynamic WEP (802.1x):** Automatically obtains up-to-date WEP keys from a key management service. This is a relatively uncommon, vendor-specific authentication mechanism.
- **WPA & WPA2 Personal:** Enables you to enter a passphrase that is used to authenticate yourself to the wireless network. Wi-Fi Protected Access (WPA and WPA2) provides improved security over the older, and simpler, WEP security. WPA Personal mode, also known as preshared key mode (PSK) is designed for home and small office networks that don't need a full-blown 802.1X authentication server—the passphrase is stored in the wireless access point.
- **WPA & WPA2 Enterprise:** Enables you to provide a username and password or identify user and certificate authority (CA) certificates and a key that was generated by the administrator of your wireless network. Certificates and keys are stored in files on your system and are used to authenticate you to an authentication server to provide access to the wireless access point. If you are using EAP or Protected EAP authentication, you will need to get a valid username and password combination from the administrator for the RADIUS server that you are accessing. WPA Enterprise mode supports additional EAP (Extensible Authentication Protocol) security mechanisms that have been certified by the Wi-Fi alliance, and can also be used with the popular, free FreeRADIUS server (`http://freeradius.org`).

Figure 13-8 shows the fields that are associated with entering WPA/WPA2 Enterprise security information using TLS (Transport Layer Security) authentication information. Clicking the fields that initially say None displays a browse dialog that enables you to locate and identify the files that contain this information.

After selecting a security mechanism and providing the information that it requires, click the IPv4 Settings tab to display the dialog that was shown earlier in Figure 25-5. Follow the instructions earlier in this chapter if you want to specify static IP address for your wireless host, which is an extremely rare circumstance. Wireless devices that use static IP addresses for security and monitoring purposes still typically retrieve them from a DHCP server that has been configured to deliver specific IP addresses to specific MAC addresses. See Bonus Chapter 10 for information about DHCP server configuration.

FIGURE 25-8

Entering WPA & WPA2 Enterprise security information



After entering all of the information required to access your wireless network, click OK to save any changes that you have made to the configuration of your wireless network connection, or click Cancel to discard those changes.

TIP **If configuration data such as wireless key information isn't being saved for a connection that is identified as a System setting, try deselecting and reselecting the Connect automatically option. A little voodoo never hurts….**

After changing wireless network settings, the Network Manager will attempt to connect to any wireless connections that you have set as Configure automatically. You can also initiate a wireless connection by left-clicking on the Network Manager applet and selecting the Connect to Hidden Wireless Network menu command. This displays a dialog from which you can select the network that you want to connect to and click Connect to attempt to initiate a connection.

TIP **As mentioned earlier, if you are having problems getting your wireless connections to work using the Network Manager, see Chapter 26 for information about command-line tools to help with debugging, or for information about installing and using the `Wicd` package, which is a great replacement for the Network Manager that often does a better job of handling some types of wireless authentication.**

Configuring a VPN Connection

Virtual private networks (VPNs) are essentially a mechanism for tunneling between two networks, enabling you to access a remote network as though it were local to your system (modulo some slight decrease in access speed). VPNs are most commonly used when connecting to business or academic networks from home or laptop computer systems, making the laptop or home system a member of a business or academic network for the duration of the VPN connection, but they can also be used as a central corporate or academic resource by permanently connecting multiple locations as parts of a single network. Prior to the development of VPNs, institutions that needed multiple locations to interact on a single network did so by using dedicated connections, referred to as leased lines, that were provided by a telecommunication provider. VPNs provide an equivalent solution that is more cost-effective and flexible, and does not require network administration by a third party such as a telecommunications carrier. VPN connections can be made from any location over an IP network, do not require purchasing or leasing dedicated network connections, and require no separate administration other than correctly configuring the VPN servers that remote clients contact in order to join the network.

As you might hope, Ubuntu's Network Manager applet makes it easy to set up VPN connections, using essentially the same mechanism that it provides for creating and editing wired and wireless network connections. The Network Manager supports OpenVPN, PPTP (Point-to-Point Tunneling Protocol), and VPNC (Cisco's VPN Concentrator) VPNs, each of which has its own configuration requirements.

NOTE **While the Network Manager applet does indeed make it easy to set up VPN connections, its support for actually connecting to them is broken at the time that this book was written. This worked in previous versions of Ubuntu and will obviously be fixed at some point, but if it's still broken when you read this, see the sidebar entitled "Setting Up and Using OpenVPN," later in this chapter, for information on establishing a VPN connection to a remote network using OpenVPN, regardless of the state of the Network Manager.**

Before starting to configure a VPN using the Network Manager applet, do the following:

- Obtain the necessary security information from the administrator of the network that you want to access via a VPN. This security information can consist of certificate files, a key file and associated password, a username and password, or some combination of these.
- Check with the administrator of the network that you want to access to determine the type of VPN that they are using at the remote end. Different VPNs require different Network Manager plug-ins.
- Check with the administrator of the network that you want to access to verify how VPN clients obtain IP address, DNS, and domain information. The VPN server usually provides these automatically when you establish a connection, but it's a good idea to check. It is inconvenient at best to be sitting in a hotel room in East Nowhere and discovering that you need some missing bit of critical information such as the IP address of a DNS server for the remote network in order to actually use your VPN connection.
- Install the Network Manager plug-in for the type of VPN that you are connecting to, using your favorite package management tool from those discussed in Chapter 23. Supported VPN frameworks and the packages that contain their associated plug-ins are OpenVPN (`network-manager-openvpn`), PPTP (`network-manager-pptp`), and VPNC (`network-manager-vpnc`).

TIP **If you are configuring a VPN connection on a laptop, it's a good idea to install all three of the VPN plug-ins. Though you'll probably only use one of these initially, it's convenient to already have them on your laptop if you are traveling and need to connect to a site that uses a different type of VPN software.**

After completing these steps, you can begin configuring your end of the VPN by doing one of the following:

- Right-click the Network manager applet in the Notification Area, select the VPN Connections menu, and select the Configure VPN menu item.
- Left-click the Network Manager applet, and select Edit Connections to display the Network Connections dialog shown in Figure 25-2. Click the VPN tab.

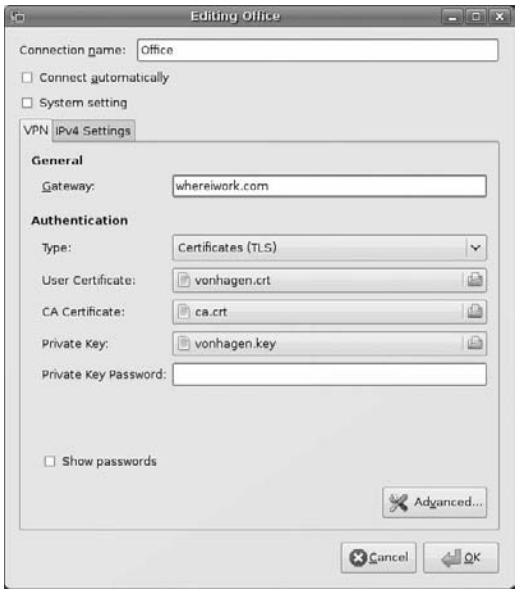
Doing either of these displays the Network Connections dialog's VPN pane. If the Add, Import, and Export buttons are not enabled, you neglected to install at least one of the plug-ins that the Network Manager uses to provide VPN access. See the final bullet in the "Before starting…" list earlier in this section.

If the system administrator of the network to which you are connecting supplied you with a configuration file for VPN clients, you can click Import to display a dialog that enables you to browse for and open that configuration file. This will create a VPN connection for you that matches the settings required by the remote server, although you may still have to supply information that may not be in a generic configuration file, such as the location of your certificates and key, your username and password, and so on. You can then select the entry for that connection in the connection list and click Edit to customize it and supply any missing information.

If you do not have a configuration file to start with and therefore need to create a completely new VPN connection, click Add. A dialog appears that provides a drop-down list from which you can select the type of VPN connection that you want to create. Select the appropriate type of VPN connection, and click Create. A dialog like the one shown in Figure 25-9 appears. The contents of this dialog will differ based on the type of VPN connection that you are defining—the example in this section uses an OpenVPN connection as an example.

FIGURE 25-9

Defining a connection to an OpenVPN VPN



First, you'll probably want to change the name of your new connection to something more evocative than "VPN Connection *N*," which is the default name format for new VPN connections. I typically use names such as "work" or "*company*-consulting" that identify the target of the VPN. But you can use any name that you want.

Next, all types of VPN connections require that you identify a Gateway, which is the hostname (or IP address) of the system that is running the VPN server software and through which you connect to the remote network.

Finally, you probably want to make sure that neither the Connect automatically or Systems setting options are selected. You probably only want to connect to a VPN under certain circumstances, and you probably don't want other users to be able to see or attempt to use any VPN connections that you have configured.

After configuring these standard items, you need to supply the authentication information that is required to establish a connection to the remote network. This information differs for each type of VPN connection, and also varies for the different authentication mechanisms that some VPN connections, such as OpenVPN connections, support. For OpenVPN connections, select the appropriate type of authentication from the Type drop-down, and then click to browse for the location of the files containing your certificates and key, enter password information if required, and so on.

After completing this dialog, you may want to click the IPv4 tab to verify that the mechanism by which your connection will get mandatory IP information matches the requirements of the remote system.

FIGURE 25-10

Verifying IP configuration for a connection to an OpenVPN VPN



The default mechanism for obtaining IP information is Automatic (VPN), which means that the VPN server acts as a standard DHCP server, providing an IP address, the identity of at least one DNS server for the remote network, the domain name for hosts on the remote network, and so on. Some VPNs are configured to only provide an IP address, and require that you identify a DNS server and supply the domain name for hosts on the remote network. Others only accept connections from specific IP addresses, which must therefore be manually configured.

Once you have supplied all of the information that is required for the type of VPN that you are connecting to, click OK to save your new connection. You can then test it by right-clicking on the Network Manager applet, selecting the VPN Connections submenu, and selecting the entry for the VPN connection that you just configured from the list at the top of the submenu that appears. Depending on the type of VPN to which you are connecting and the information that you supplied in the VPN dialog like the one shown in Figure 25-9, you may be prompted for a username, password, or both in order to establish your connection. The icon for the Network manager applet will change to show a swirling globe and arrows as your system establishes your connection, and a status message will be displayed when the connection attempt completes.

If you get an error or connection failure message, edit and verify the settings for your connection and try connecting again. If you are still having problems connecting, your best bet is to search the Web using a query like "Ubuntu Network Manager VPN problem" to see if there are any known problems with the type of connection that you are trying to establish. You can also skip using the Network Manager and try to establish your VPN connection manually. The client-side server that your system must run in order to connect to a specific type of VPN was installed along with the Network Manager plug-in for that type of VPN, and is usually integrated into the system boot process. However, connections will not be automatically started if you need to supply a password or other information during the connection process. See the sidebar "Setting Up and Using OpenVPN" for an example of manually creating configuration files and establishing a connection to a remote OpenVPN VPN server.

Setting Up and Using OpenVPN

OpenVPN is a free, open source VPN package that is released under the GPL, and is one of the most popular VPN packages in use today. OpenVPN uses the OpenSSL library and SSL/TLS (Secure Sockets Layer/Transport Layer Security) for secure, encrypted communications, and uses preshared secret keys, certificates, or username/password combinations for authentication. Beyond its popularity on Linux systems, OpenVPN is also available for operating systems, including FreeBSD, Mac OS X, NetBSD, OpenBSD, Solaris, and Microsoft Windows 2000, XP, Vista, and 2008 Server.

Installing OpenVPN on your system can be done manually by installing the `openvpn` package, or automatically as a requirement for installing the `network-manager-openvpn` package that provides the Network Manager's support for creating and managing OpenVPN connections. As part of installing OpenVPN, it is integrated into your Ubuntu or Kubuntu system's startup process. As part of its startup process, OpenVPN will try to automatically initiate any VPN connections for which configuration files are present in the directory `/etc/openvpn`. These configuration files can have any name, but must have the file extension `.conf`. Each configuration file identifies the OpenVPN parameters that are required to establish a VPN connection to a remote network, including the name and location of any certificate and key files that each connection requires. However, any connections that require specifying a password to unlock or verify a private key will fail during the boot process, because OpenVPN does not support storing username and password information in configuration files unless it has been specially compiled to do so. Storing username and password information is usually not supported because storing this information in clear text in a configuration file is a huge potential security hole.

continued

continued

The following is a sample OpenVPN configuration file:

```
port 1194
dev tun
remote whereiwork.com
tls-client
ca ca.crt
cert vonhagen.crt
key vonhagen.key
mtu-test
tun-mtu 1500
tun-mtu-extra 32
mssfix 1450
pull
comp-lzo
verb 0
```

These `remote`, `ca`, `cert`, and `key` entries in OpenVPN configuration files will be different for each VPN connection. The `remote` keyword identifies the host that is running the OpenVPN server for a connection, the `ca` keyword identifies the file containing the root Certificate Authority certificate for the remote site, the `cert` keyword identifies the file containing the client's certificate, and the `key` keyword identifies the file containing the private key for the client. In this example, the latter two are user-specific, but they can also be sitewide to enable sitewide connections. User-specific key files are typically owned by the root user and are stored as mode 600 so that they can only be read by the root user.

The "key" issue in whether an OpenVPN connection can be automatically created at boot time is whether the private key file requires a password in order for it to be used. System administrators can create key files that do not require passwords using the `-passin pass:`*`password`* option (where *`password`* is a password that will be used automatically) when creating a key file using the `openssl` utility. Such key files can be used to automatically set up "always-on" VPN connections at boot time. VPN connections that use a key file that requires a password and for which all files are stored in `/etc/openvpn` can be initiated at any time by executing the following command from any command-line environment:

```
sudo /etc/init.d/openvpn restart
```

This command will temporarily stop all current OpenVPN connections and will then attempt to start all OpenVPN connections for which configuration files exist in `/etc/openvpn`, prompting you for any passwords that are required. A more elegant way of starting an individual OpenVPN connection without interrupting any existing connections is to execute a command like the following from any command-line environment:

```
sudo openvpn --cd /etc/openvpn --daemon --config configfile.conf
```

The `--cd` option specifies the directory in which any files without path names (such as those in the previous example) can be found, the `--daemon` option tells the `openvpn` daemon to detach from the command-line environment from which it was started if the daemon can start correctly, and the `--config` option specifies the name of the configuration file for the OpenVPN connection that you want to start. You will be prompted for the password for the key file identified in the configuration file, and the OpenVPN daemon will create the connection and put itself in the background in daemon mode. As you can see, this command also enables you

continued

continued

to store user-specific OpenVPN configuration files in any directory by changing the value of the `--cd` option to the name of that directory.

OpenVPN connections that you started will persist after you log out, so you should terminate any user-specific OpenVPN connections by using the `ps` command to identify the process ID for the OpenVPN server that you started and using the `kill` command to kill that process, which will also remove any routing table entries for that OpenVPN connection. If you are only running a single instance of the `openvpn` process, you can simplify this by using a command like `kill -9 `pidof openvpn`` to shoot it down.

Using the Network Manager on Kubuntu

As discussed in the introduction to the "Configuring Your Network Hardware" section of this chapter, both Ubuntu and Kubuntu systems install and use the Network Manager to locate, identify, and configure your network hardware. The previous section discussed the Network Manager applet that is used as a front end for the Network Manager to configure and manage both wired and wireless network connections on GNOME-based Ubuntu desktop systems. As a GNOME panel applet, this utility is not installed by default on Kubuntu systems, which instead use a front end to the Network Manager known as KNetworkManager that runs in the KDE System Tray. The next few sections explain how to use KNetworkManager to examine and customize wired, wireless, and VPN connections on Kubuntu systems. Both the KDE 4-based Kubuntu 8.10 and newer and the traditional KDE 3-based Kubuntu 8.04 and earlier systems use KNetworkManager to create and manage their network connections. There are relatively few differences between using KNetworkManager on these types of systems. The majority of the screenshots in this section were, therefore, taken on a Kubuntu 8.04 system, but I'll provide KDE 4 screenshots whenever there is any substantial difference in the menus and dialogs in the two versions of KDE.

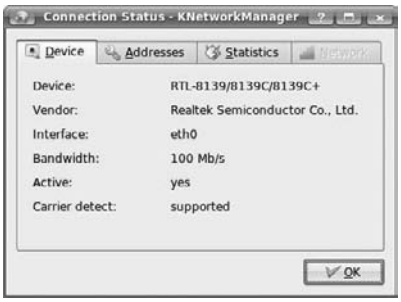
By default, all interaction with network devices on Kubuntu systems is done through the KNetworkManager that runs in the KDE System Tray, located by default in the lower-right corner of the kicker panel. This applet uses different icons on KDE 3 and KDE 4 systems, which also change depending on the type of your current network connection and its status:

- **Kubuntu 8.04 and earlier (KDE 3):** If you are connected via a wired connection, the icon looks like an RJ-45 network jack (or an ice cream bar to some of my less network-savvy friends). If you are not connected to a network, the icon shows both the socket for an RJ-45 jack and the jack itself, and also displays a red "x" to highlight the fact that you are not currently connected. If you are only connected via a wireless connection, the icon looks like a blue histogram. A rotating gear replaces the icon whenever any of your network connections is changing state.
- **Kubuntu 8.10 and later (KDE 4):** If you are connected via a wired connection, the icon displays a green globe that is evocative of the earth. If you are not connected to a network, the icon shows a gray globe. If you are only connected via a wireless connection, the icon looks like a blue histogram. A rotating gear replaces the icon whenever any of your network connections is changing state.

On a Kubuntu 8.04 or earlier system, left-clicking on whatever KNetworkManager icon you see displays a dialog that shows the connection status information for the Ethernet device that is currently active, as shown in Figure 25-11. (On Kubuntu 8.10 systems, left-clicking doesn't display this dialog, but simply does the same thing as right-clicking.)

FIGURE 25-11

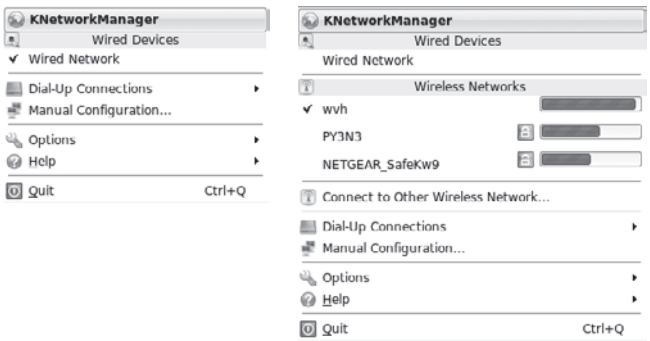
Connection status information from the KDE 3 KNetworkManager applet



Right-clicking whatever KNetworkManager icon, you see displays menus like the one shown in Figure 25-12 that list the current status of all of the network connections that were detected on your system, enable you to connect to other wireless networks, enable you to manually configure any available network device, and provide submenus for configuring dial-up connections, setting networking and KNetworkManager options, and getting help.

FIGURE 25-12

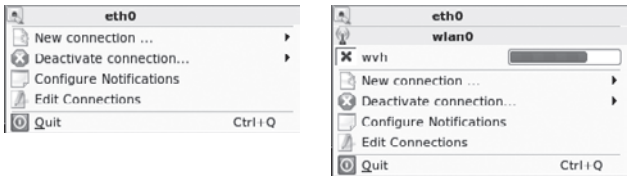
General status information and other options from the KDE 3 KNetworkManager



The left portion of Figure 25-12 shows the KNetworkManager menu on a system that has a single wired Ethernet interface. The checkmark to the left of the name of that interface shows that it is currently active. The right portion of Figure 25-12 shows the KNetworkManager menu on a system that has a single wired Ethernet interface, and which also has a wireless Ethernet interface that has detected the availability of three wireless networks. The wired Ethernet interface is not checked, indicating that it is not in use, and the checkmark to the right of the wireless network `wvh` shows that the system's wireless Ethernet interface is currently connected to that wireless network. The bars to the right of each wireless network indicate the strength of the signal from an access point for that wireless network. The icon to the left of two of those bars indicates that these are secure networks that require authentication in order to access them, which the network named `wvh` does not. Figure 25-13 shows the pop-up menus that you see when right-clicking on the KNetworkManager tray icon on a KDE 4-based Kubuntu 8.10 or better system.

FIGURE 25-13

General status information and other options from the KDE 4 KNetworkManager



As you can see, the pop-up menus for the KDE 3 and KDE 4 KNetworkManager are somewhat different. The KDE 4 KNetworkManager menus don't provide menu groups that explicitly identify which of your interfaces are wired and wireless—you have to differentiate between them based on the icons that are displayed to the left of the device names. The ability to connect to another wireless network is now done by selecting wireless from the New connection submenu. The ability to manually configure an existing network interface is now done by selecting the Edit Connections menu item, and selecting the interface that you want to reconfigure in the Edit Connections dialog, which is discussed later in this section. The ability to configure the notifications that your system displays in response to changing network state is now a separate menu entry named Configure Notifications, which you can choose, rather than having to access it from the Network Settings dialog that was displayed by selecting Options ➪ Configure from the KDE 3 KNetworkManager pop-up menu.

To configure your network connections on Kubuntu 8.04 and earlier systems, right-click on the KNetworkManager tray icon and select Manual Configuration to display the dialog shown in Figure 25-14. On a KDE 3-based Kubuntu system, you can also select the K Menu ➪ System Settings menu item and select Network Settings to display KDE's network configuration applet. However, the KDE 4-based System Settings application doesn't provide similar functionality, so I'd suggest using the KNetworkManager approach if you're using KNetworkManager in the first place.

FIGURE 25-14

The KDE 3 KNetworkManager's Manual Configuration dialog

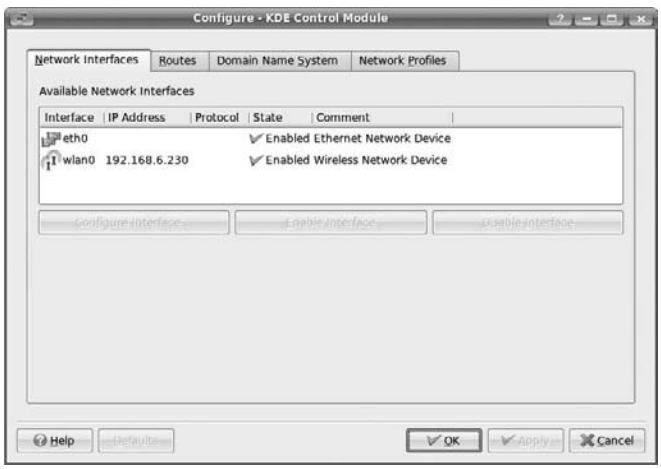
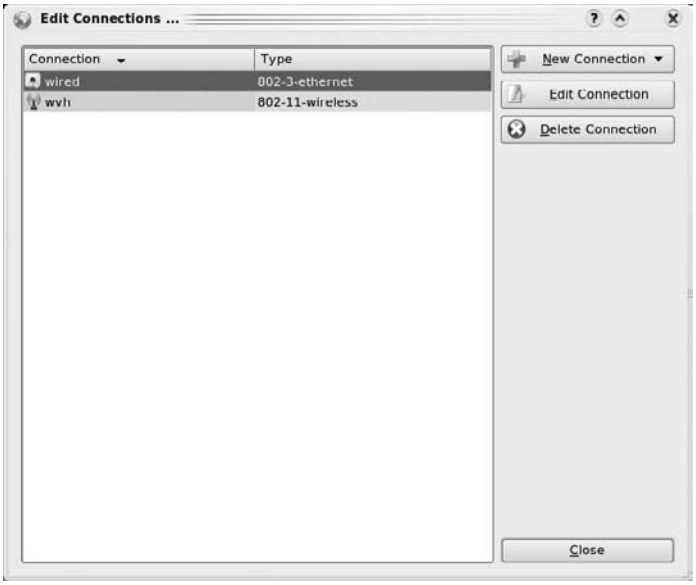


FIGURE 25-15

The KDE 4 KNetworkManager's Manual Configuration dialog



On Kubuntu 8.10 and later systems, right-click on the KNetworkManager tray icon and select the Edit Connections menu item to display a dialog like the one shown in Figure 25-15.

The next two sections explain how to configure wired and wireless connections using the KNetworkManager applet on KDE 3-based and KDE 4-based Kubuntu systems. These sections group wired and wireless configuration together, organizing them by KDE version because the KDE 3 and KDE 4 dialogs and techniques for network configuration are so different.

Remember that the Network Manager is designed to automate enabling and configuring the network interfaces on your system. In most cases, you won't have to actually do anything for standard wired and wireless Ethernet connections—if you turn on your Kubuntu system and its network interface(s) come up automatically and correctly, you don't have to do any additional configuration and can just start reading mail, surfing the Web, or any other activity that requires a network connection. However, if your interfaces didn't start automatically, used DHCP where you want to use a static IP address, if you want to connect to a different wireless network, or if you want to create a connection to a virtual private network, read on.

TIP **As mentioned earlier in this chapter, if you are having problems getting your wired or wireless connections to work using the KNetworkManager, see Chapter 26 for information about command-line tools to help with debugging, or for information about installing and using the `Wicd` package, which is a great replacement for the Network Manager that often does a better job of handling some types of wireless authentication.**

Configuring Wired and Wireless Connections on KDE 3-Based Systems

To examine or modify how a specific wired connection is configured, select your wired Ethernet interface from the dialog shown in Figure 25-14, and click Configure Interface. This displays a dialog like the one shown in Figure 25-16.

FIGURE 25-16

Configuring a wired network connection, KDE 3 style



To examine or modify how a specific wireless connection is configured, select your wireless Ethernet interface from the dialog shown in Figure 25-14, and click Configure Interface. This displays a dialog like the one shown in Figure 25-17.

FIGURE 25-17

Configuring a wireless network connection, KDE 3 style

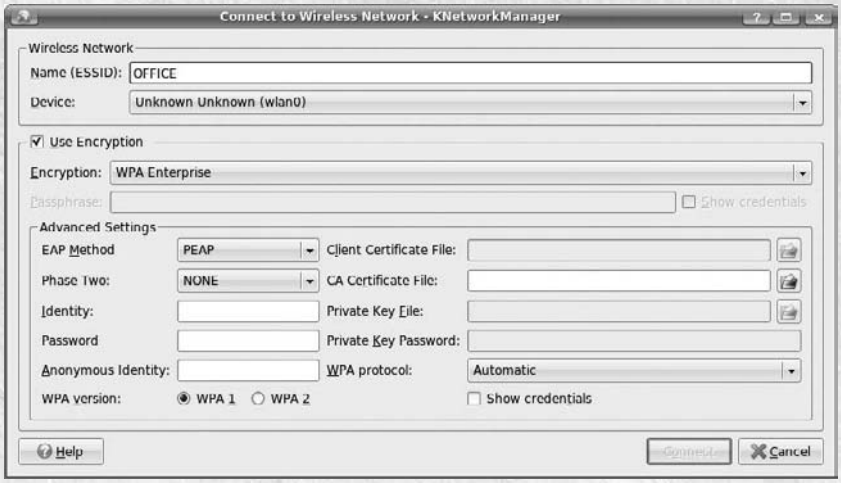


If you want this interface to use DHCP, select the Automatic radio button. If you want this interface to have a static IP address, click the Manual radio button and enter the IP address and Netmask that you want to use. You can also click the Advanced Setting button to display additional fields that enable you to specify Broadcast and Gateway addresses for this interface, though the default broadcast address for your network (XXX.XXX.XXX.254) is usually sufficient, and the gateway address is usually specified on the Routes tab of the dialog shown in Figure 25-14, which is shown in Figure 25-18. To enable this interface whenever you restart your computer system, select the Activate when the computer starts checkbox.

If you are configuring a wireless interface, you should specify the ESSID of the access point that you want to connect to, and the WEP key for "secure" networks that use WEP keys to enable access.

Wireless Security Beyond WEP on KDE 3–Based Systems

If you want to connect to a wireless network that uses more meaningful security than WEP, you do not have to reconfigure your wireless interface. You can do this by right-clicking on the KNetworkManager icon and selecting the Connect to Other Wireless Network menu command. This displays the following:



This dialog enables you to specify another ESSID and select the Use Encryption option to enable you to select between WEP Passphrase, WEP 40/128-bit hex and ASCII keys, WPA Personal, and WPA Enterprise encryption. (See the section entitled "Configuring Wireless Connections on Ubuntu" for detailed information about each of these types of wireless security.)

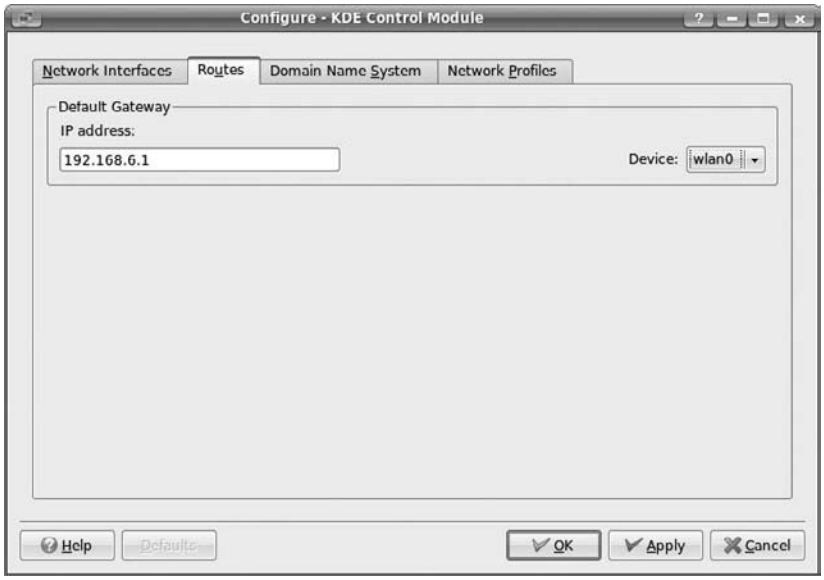
The figure in this sidebar shows the settings for using WPA Enterprise security. To identify the location of any of the files that are required for WPA security, click the file folder icon to display a Browse dialog that enables you to navigate to and select the appropriate file on your system.

After making any changes that are required by the type of security used by the wireless network that you want to access, click Connect to establish a connection to the newly defined wireless network using the associated security model.

After making any changes that you want to make, click OK to save your changes, close this dialog, and return to the dialog shown in Figure 25-14. Click the Routes tab to display the pane shown in Figure 25-18.

FIGURE 25-18

Configuring your default gateway, KDE 3 style



Enter the IP address of your default gateway in the IP address field, and select the Ethernet interface that is associated with this gateway from the Device drop-down list. Click Apply to save your changes, and select the Domain Name System tab to display a pane like the one shown in Figure 25-19.

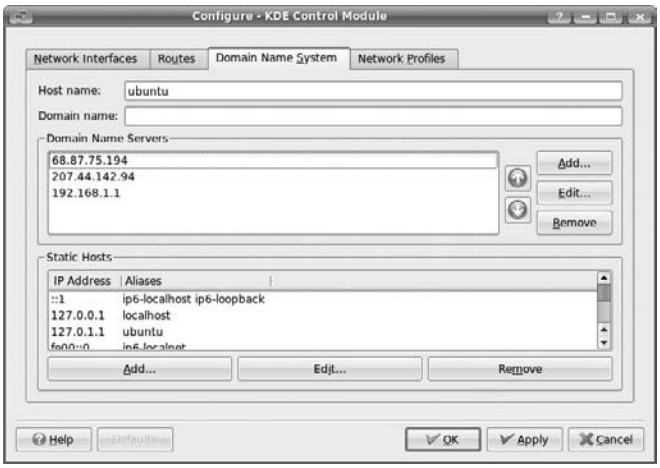
This pane displays the name of your host, the Domain Name Servers (DNS) that you have previously defined (or which were retrieved from your DHCP server), and any static hosts that are predefined in your `/etc/hosts` file. If you are using a static IP address with the interface that you are configuring, you will want to make sure that you have defined at least one DNS server. To add a new DNS server, click add, enter its IP address in the dialog that appears, and click Add to close that dialog and add it to the list shown in Figure 25-19.

NOTE **Because DNS servers are the usual source of information that map IP addresses to hostnames, you should enter only IP addresses in this dialog. If you specify a hostname, your system will need to use a DNS server to figure out the IP address associated with that DNS server, which would cause a nasty chicken-and-egg loop.**

Once you're finished verifying or specifying DNS servers, you can click OK to save your changes and close the Manual Configuration dialog.

VFIGURE 25-19

Configuring DNS servers, KDE 3 style



To activate a newly configured Ethernet interface, right-click on the KNetworkManager icon and select the interface that you want to activate. The KNetworkManager icon will be replaced with a spinning gear icon as the interface is reconfigured, and the appropriate icon for the type of interface that you activated will be displayed once that interface is fully configured and active.

Configuring Wired and Wireless Connections on KDE 4-Based Systems

To examine or modify how a specific wired connection is configured, select your wired Ethernet interface from the dialog shown in Figure 25-15, and click Configure Interface. This displays a dialog like the one shown in Figure 25-20.

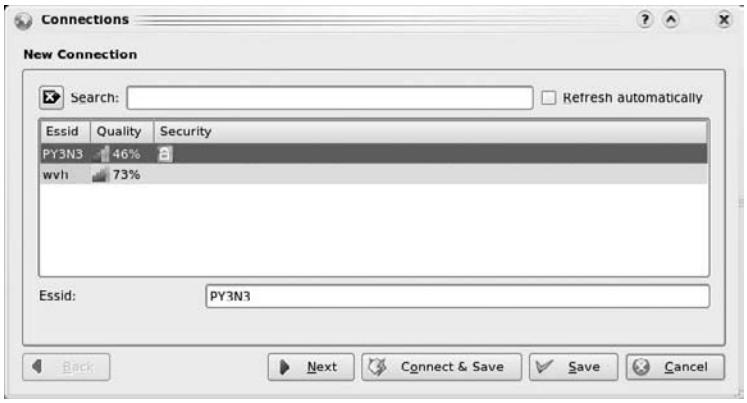
FIGURE 25-20

Configuring a wired network connection, KDE 4 style



FIGURE 25-21

Configuring a wireless network connection, KDE 4 style



To examine or modify how a specific wireless connection is configured, select your wireless Ethernet interface from the dialog shown in Figure 25-15, and click Configure Interface. This displays a dialog like the one shown in Figure 25-21.

TIP **To quickly create a new wired or wireless connection, you can right-click the KNetwork-Manager icon, select the New connection menu item, and select the interface for which you want to create a new configuration entry. This displays the dialogs shown in Figures 25-20 and 25-21 without displaying the intermediate Edit Connections dialog.**

If you are configuring a wired interface and want this interface to have a static IP address, click the Use manual IP configuration checkbox and enter the IP address, Netmask, Gateway, DNS Addresses, and DNS Search (domain name) that you want to use for this connection. Click next to proceed to the dialog shown in Figure 25-22.

FIGURE 25-22

Completing static IP address Information for a wired connection, KDE 4 style



Enter a name for the new wired interface definition, and select the Autoconnect checkbox if you want this interface to be automatically enabled and configured the next time that you start your system. Finally, click Connect & Save to save the new entry and begin using it, or click Save to save the new interface definition without switching to it.

If you are configuring a wireless interface, select the ESSID of the wireless network that you want to connect to, and click next to proceed to the dialog shown in Figure 25-23.

FIGURE 25-23

Specifying wireless security, KDE 4 style



To finish configuring a wireless interface, select the type of security (WEP or IEEE 802.1X) that you want to use from the drop-down Type list. (See the section entitled "Configuring Wireless Connections on Ubuntu" for detailed information about these types of wireless security.) The fields required for the type of security that you have selected are displayed at the bottom of the dialog. Enter values in those fields and click Connect & Save to save your changes and connect to the wireless network associated with your new configuration, or click Save to save the newly defined entry without switching to it.

NOTE **The version of KNetworkManager that was available when this book was updated only supports a fairly minimal set of wireless authentication and encryption protocols. If this is still the case when you read this, see the discussion of `Wicd` in Chapter 26 for information about an alternative to the Network Manager (and KNetworkManager) that provides superior support for a wide range of wireless authentication and encryption protocols.**

To activate a newly configured Ethernet interface, right-click on the KNetworkManager icon and select the interface that you want to activate. The KNetworkManager icon will be replaced with a spinning gear icon as the interface is reconfigured, and the appropriate icon for the type of interface that you activated will be displayed once that interface is fully configured and active.

Configuration Files for Network Interfaces

Tools such as Network Manager, discussed in the previous two sections, and `Wicd`, discussed in Chapter 26, are great if you want flexible, dynamically configured Ethernet interfaces on your system. But what if these don't work for you or if (gasp!) your system doesn't go anywhere or only uses one or two configurations? In these cases, you can eliminate the overhead that mechanisms such as Network Manager and `Wicd` introduce by falling back to the original network configuration mechanism that was pioneered by the Debian Linux distribution and has therefore been inherited by its offspring such as Ubuntu and Kubuntu systems: text files.

Network interfaces that rarely change can easily be configured by creating entries for them in the file `/etc/network/interfaces`, which is the core network configuration file used on Ubuntu and Kubuntu systems. This file supports both wired and wireless interfaces, including security configuration for wireless interfaces. In order to configure your network interfaces using this file, you must first make sure that neither the Network Manager nor `Wicd` are running when you start your system. To remove either of these network configuration mechanisms from your system's startup sequence, you must remove the symbolic links `S28NetworkManager` or `S20wicd` from the directory `/etc/rc2.d`, which is the directory that contains the files that bring a Ubuntu or Kubuntu system up to its default run level, run level 2. You can simply remove these links, but I typically like to remember what I've enabled and disabled on my systems. I, therefore, create a directory called `/etc/rc2.d/DISABLED` and move whichever of these symbolic links is present on my system into that directory. This will prevent those network configuration mechanisms from running the next time that you start you system.

Once you've disabled automatic network configuration mechanisms for your next reboot, you can create entries for your network interfaces in the file `/etc/network/interfaces` using your favorite text editor. This file will already contain an entry for the loopback interface on your system. The easiest type of entry in this file is an entry for an Ethernet interface that uses DHCP to obtain its IP address and other network configuration information. An entry that defines the Ethernet interface `eth0` as one that uses DHCP will look like the following:

```
auto eth0
iface eth0 inet dhcp
```

The `auto` line in this entry states that the interface `eth0` will be automatically enabled when the system starts. The `iface` line in this entry identifies the beginning of the interface definition for `eth0`, which is a TCP/IP (`inet`) device that uses `dhcp` to obtain its configuration information.

Entries that assign static IP addresses to specific interfaces are equally simple. For example, an entry that assigns a static IP address for the Ethernet interface `eth0` will look like the following:

```
auto eth0
iface eth0 inet static
     address 192.168.6.64
     netmask 255.255.255.0
     gateway 192.168.6.1
```

The `auto` line in this entry states that the interface `eth0` will be automatically enabled when the system starts. The `iface` line in this entry identifies the beginning of the interface definition for `eth0`, which is a TCP/IP (`inet`) device with `static` configuration. The next three lines define the IP `address`, `netmask`, and `gateway` for that Ethernet interface.

Manual DNS Configuration

If your network interfaces use static IP addresses, your system will not obtain general network configuration information such as DNS server addresses from a DHCP server. You should, therefore, manually create and populate the file `/etc/resolv.conf` that your system uses to identify its name servers. A sample `/etc/resolv.conf` file that uses open DNS servers from OpenDNS.org is the following:

```
domain vonhagen.org
nameserver 208.67.222.222
nameserver 208.67.220.220
```

The first line in this file identifies the default domain that your system should search for hosts that are specified without domain names. (You will probably use a default domain that is something other than my home domain.) You should add entries for any DNS servers that are provided by that default domain, which should appear before entries for the DNS servers from open services such as OpenDNS. When your system tries to find the IP address associated with a specific host, it will query the DNS servers in the order that they are listed in this file.

For information about other DNS servers that are available to the general public, see `www.dnsserverlist.org`.

Wireless Ethernet interfaces are similarly easy to define in your `/etc/network/interfaces` file. For example, an entry that defines the Ethernet interface `eth1` as one that uses DHCP and the wireless network with the ESSID "wvh" will look like the following:

```
auto eth1
iface eth1 inet dhcp
    wireless-essid wvh
```

The `auto` line in this entry states that the interface `eth1` will be automatically enabled when the system starts. The `iface` line in this entry identifies the beginning of the interface definition for `eth1`, which is a TCP/IP (`inet`) device that uses `dhcp` to obtain its configuration information. The `wireless-essid` entry defines the name of the wireless network that this entry should connect to. You can use the keyword `any` instead of identifying a specific `wireless-essid` if you want to connect to any open wireless network that is available.

Entries such as this one are fine for open wireless networks. Entries associated with WEP wireless networks require at least one additional entry to specify the wireless key, as in the following example:

```
auto eth1
iface eth1 inet dhcp
    wireless-essid wvh
    wireless-key 0102030405
```

If the wireless network that you want to connect to uses WPA authentication, you will have to install and configure a package such as `wpasupplicant`, and integrate that package into your wireless configuration

entries. Describing how to install and configure this package is (unfortunately) outside the scope of this section. See the forum postings at `http://ubuntuforums.org/showthread.php?t=202834` for more information about using wireless security mechanisms that are more powerful than WEP.

You can also define static IP addresses for wireless interfaces, as in the following example:

```
auto eth1
iface eth1 inet static
    address 192.168.6.42
    netmask 255.255.255.0
    gateway 192.168.6.1
    wireless-essid wvh
    wireless-key 0102030405
```

The file `/etc/network/interfaces` provides a simple, low-overhead mechanism for configuring network interfaces that are either fairly simple or don't change very often. The examples in this section have only scratched the surface of possible entries in this file, but have identified the most important ones. For complete information about the possible entries in this file for wired Ethernet interfaces, see the online man page for this file by typing the command `man 5 interfaces` from any command-line interface. The web is the best resource for locating complete information about the keywords in this file that are related to wireless network configuration.

Configuring Modems and Dial-Up Connections

Though broadband Internet access is common nowadays, dial-up connections using protocols such as PPP (Peer-to-Peer Protocol) are still the way in which some people connect to the Internet. Depending on your geographic location, this may be your only option for connecting to the Internet. Unfortunately, many modern modems are so-called "Win-modems," which saves hardware companies significant money by using Microsoft Windows device drivers and the operating system itself to do the MOdulation/DEModulation that gives modems their name. Unfortunately, such drivers are rarely provided for Linux. See `http://linmodems.org` for information on whether your modem is a Win-modem and, if so, software that may be able to make it work with Linux.

Using a modem to connect to the Net is rare nowadays and, for that reason, Ubuntu and Kubuntu systems don't install packages for modem configuration and connection setup by default. If you need to use a modem and have one that works on Linux systems, the graphical `gnome-ppp` and `kppp` packages can make it easy for you to configure and initiate modem-based network connections. Similarly, the `wvdial` package that is available in the Ubuntu repositories is a great command-line utility application to use to initiate a PPP or SLIP connection over your modem. See the man page for `wvdial`, its FAQ (`http://alumnit.ca/wiki/index.php?page=WvDialFAQ`), or its home page (`http://freshmeat.net/projects/wvdial`) for information on configuring and using this package. If you're reading this, I apologize for not discussing these packages in detail, but I don't have a dial-up account anymore (or enough phone lines to test with the stack of retired modems that are in my basement).

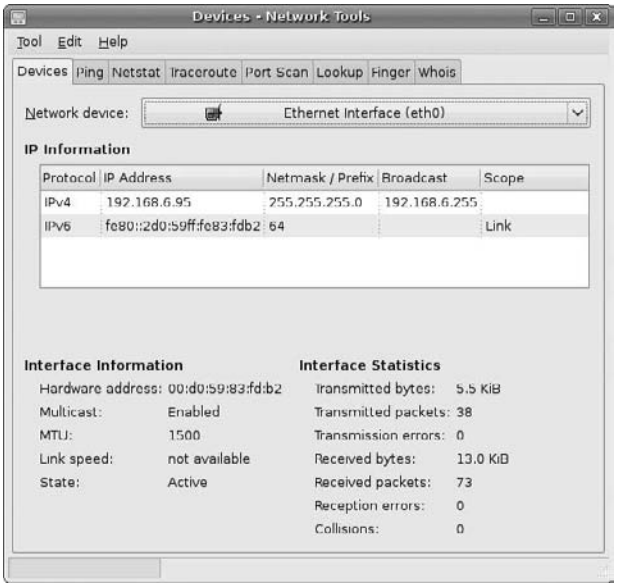
Network Testing with Network Tools

To maintain its tradition of easy graphical network tools, Ubuntu provides a convenient graphical tool that simplifies examining the current configuration of any of your network interfaces. Ubuntu provides GNOME's Network Tools application to give you a graphical display of network configuration information, as well as easy graphical access to a variety of network tools. Select the System ➪ Administration ➪ Network Tools menu item to start the Network Tools application. This displays a dialog like the one shown in Figure 25-24.

NOTE **Sorry, Kubuntu fans, the `gnome-nettool` package that is discussed in this section is a GNOME application and is, therefore, only installed on Ubuntu systems. You can certainly install it on a Kubuntu system, but that will also require installing its GNOME and GTK dependencies. Personally, I don't care about disk space, but do care about convenience, so I always install this on my Kubuntu systems and typically also often install it on the Kubuntu systems that I set up for friends and relatives. On the other hand, this tool is really just a single convenient interface to a number of utilities that are already available from any command-line environment. Whether or not to install it on a Kubuntu system is completely up to you.**

FIGURE 25-24

Information about eth0 in the Network Tools application



By default, the Network Tools application shows information about your system's loopback interface. To see information about a specific interface, click the Network device drop-down menu and select the Ethernet interface that you're interested in. Figure 25-24 actually shows information about the primary physical Ethernet interface, `eth0`, on one of my test systems.

The easiest and fastest way to identify the current configuration of one of your Ethernet interfaces will probably always be to run the `ifconfig interface-name` command in an `xterm` or GNOME Terminal window, which will display output like the following:

```
$ ifconfig eth0
eth0      Link encap:Ethernet  HWaddr 00:16:EC:CA:AE:20
          inet addr:192.168.6.64  Bcast:192.168.6.255  Mask:255.255.255.0
          inet6 addr: fe80::216:ecff:feca:ae20/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:32034 errors:0 dropped:0 overruns:0 frame:0
          TX packets:28776 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:12108607 (11.5 MiB)  TX bytes:4603767 (4.3 MiB)
```

As you can see, the text display of Ethernet interface information provided by the `ifconfig` command still requires a certain amount of interpretation when compared to the friendlier display of information shown in Figure 25-24.

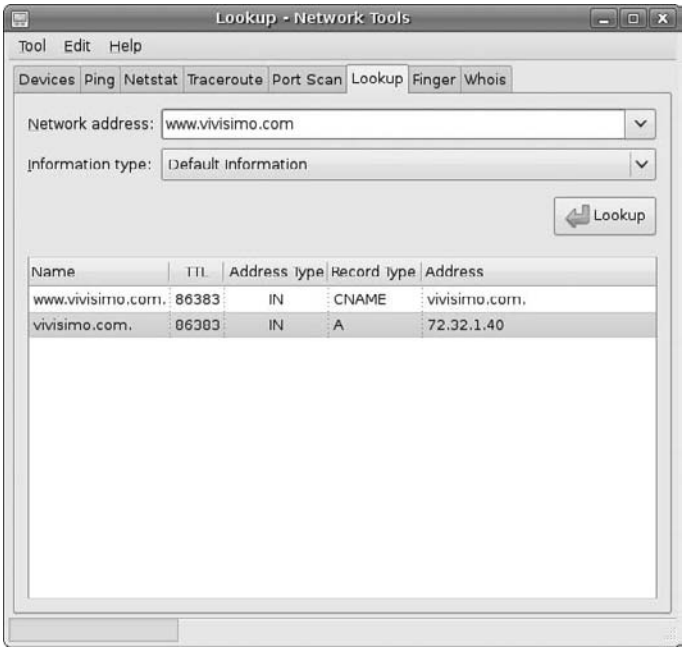
In addition to a more readable display of basic network configuration information, the Network Tools application supports the graphical display of information produced by several standard network utilities, which traditionally operate only in text mode. The tabs provided in the Network Tools application, along with the purpose of each tab, are the following from left to right:

- **Devices:** Displays configuration and traffic summary information for each available network interface on the system. This corresponds to the information provided by the traditional Linux/UNIX command-line `ifconfig` application.
- **Ping:** Displays connectivity and availability information by sending packets to a specified host or IP address, and displays elapsed time and success/failure information. This corresponds to the information provided by the traditional Linux/UNIX command-line ping application.
- **Netstat:** Displays status information about all active and available TCP and UDP network ports on the system. This corresponds to the information provided by the traditional Linux/UNIX command-line `netstat` application.
- **Traceroute:** Displays the systems through which communication to a specified host pass and the time required for each intersystem communication, known as a *hop*. This corresponds to the information provided by the traditional Linux/UNIX command-line `traceroute` application.
- **Port Scan:** Displays information about available ports and services on a specified remote machine. This roughly corresponds to the information provided by the traditional Linux/UNIX command-line `nmap` application.
- **Lookup:** Displays IP address information and available DNS aliases for a specified system. This roughly corresponds to the information provided by the traditional Linux/UNIX command-line `nslookup` or host applications.
- **Finger:** Displays any available personal information about a specific user or a specified host. This corresponds to the information provided by the traditional Linux/UNIX command-line `finger` application. Few hosts provide this information any longer.
- **Whois:** Displays information about the registrant and technical contact for a specified Internet domain. This corresponds to the information provided by the traditional Linux/UNIX command-line `whois` or `bwhois` applications.

As an example of the type of information and display provided by the Network Tools application, Figure 25-25 shows DNS hostname lookup information as shown on its Lookup tab for the host `www.vivisimo.com`.

FIGURE 25-25

Graphical display of DNS name lookup



Tips for Securing Your System

System security is an open-ended topic because it has so many different aspects. These include physical security, login authentication, file and filesystem protections, and so on. Entire books have been written about security topics, and more are doubtless on the way. As mentioned in the introduction to this chapter, security in all forms will become an ever-increasing concern because of the increasing ubiquity of networking and the increasing availability of easy-to-use tools for probing, exploring, and breaking into remote machines.

The following are some specific suggestions for increasing the security of your system on a network. As you'd expect, these include some aspects of other security topics but also have their own unique concerns:

- If you are using an off-the-shelf home gateway, change the password before you put it into service. You should also change the name of the authenticated user if possible. If I had a nickel for how many systems have been broken into because people didn't change default passwords, I wouldn't even know how many nickels I had because most of these break-ins go unnoticed.

- Disable any unnecessary services on your system. You can use the Network Tools Port Scan tab to identify ports on your system that are listening from requests for services. Disable any services that you are not using through a tool such as the Boot Up Manager, which was discussed in the section of Chapter 19 entitled "Optimizing the Ubuntu Boot Process."
- Remove accounts for any users that are no longer using your system. This includes system accounts that were created for use by or with services that you are no longer running on your system.
- Always keep your system up to date using the Ubuntu Update Manager. Patches to system and application software are released for a (good) reason.
- Monitor important system log files regularly. The `/var/log/messages` and `/var/log/syslog` files can be an important source of information about who is trying to break into your system and how.
- Change your password regularly. Ubuntu's dependence on the `sudo` command rather than the traditional root account for system administration tasks is a useful obfuscation, but your dedicated cracker in Beijing often doesn't have anything better to do than try and try again.

As mentioned previously, security is your responsibility. Some interesting applications are available to test and probe your own system, which can be both educational and useful. My long-term favorites are:

- **chrootkit:** Checks for "root kits," which is the term for precompiled sets of hacked applications that are often installed on systems that have been broken into. These root kits both make it easier for a cracker to get into your system again and also collect additional login/password information from a cracked system.
- **nmap:** Probes network connectivity on your machine and identifies potential problem points.

As you might expect, both of these applications are available in the Ubuntu repositories and can easily be installed on your system any of the package management tools discussed in Chapter 23.

Installing a Firewall

Firewall is the term used to describe a system that sits between one or more computer systems and monitors and manages network traffic. Just as with the firewall in your automobile, which prevents a fire in the engine compartment from proceeding into the passenger compartment and incinerating its occupants, a network firewall is intended to prevent malicious, spurious, or unnecessary network traffic from moving through it. Many firewalls serve multiple functions, also performing services such as Network Address Translation (NAT), but their primary purpose is to protect against network attacks and other unwelcome intrusions.

On modern Linux systems, firewalling is typically done using kernel modules that support a packet filtering framework known as `netfilter`, and an associated interface and user-space command known as `iptables`. *Packet filtering* refers to the ability to analyze network packets and perform various actions based on their source, destination, type, or other information that they contain. Because support for packet filtering is built into the Linux kernel, a Linux system that is directly connected to the Internet can serve as its own firewall, monitoring and managing network traffic before that traffic actually gets to any daemons or network-aware processes that it is running. Of course, a dedicated device or Linux system can also serve as a firewall, and many vendors sell prepackaged solutions that do just that. The fact that many of these off-the-shelf systems run Linux and use the `netfilter/iptables` mechanism to implement their firewalling solutions is just proof of the power of the Linux kernel's built-in support for packet filtering.

Whether or not a Linux system actually requires a firewall is a hot debate topic among Linux fans. Standard Ubuntu and Kubuntu Desktop installations do not expose any open ports to an outside network, so there are no network ports that need to be protected. This is not true, of course, for Ubuntu server systems that expose ports for services such as DNS, e-mail, SSH, a Web server, and so on, so a firewall is always a good idea for any server system.

TIP **If you are using your Ubuntu or Kubuntu system in an environment that is already protected by a firewall, you probably do not need to set up a firewall on your system. You should, however, make sure that the firewall that your system is located behind is actually doing the right thing by checking with the manufacturer, your IT group in a business or academic environment, or your Internet service provider. Just because a box has "Firewall" printed on it doesn't mean that it is actually doing anything.**

As far as Ubuntu and Kubuntu desktop systems go, you will probably find yourself opening up some ports on a desktop installation as you use your system over time, and a `netfilter/iptables` firewall introduces very little overhead on a desktop system, so I suggest that you always install at least a simple firewall. This way, if you subsequently increase the exposure of your system by opening ports, the firewall will already be in place. You may want to revisit your initial firewall implementation in the future, but you will at least have some protection even if you neglect firewalling in your excitement to make some new service available from your Ubuntu or Kubuntu system. Installing a simple firewall by default is also a good idea if you are setting up systems for friends, relatives, or small businesses where you may not always have complete control over what they add to or activate on their systems.

Overview of Linux Firewalling and Packet Filtering

The packet filtering mechanism used by the current Linux kernel (2.6.xx) is a combination of a loadable kernel module framework and API called netfilter, and an interface and associated and user-space administrative command called `iptables`. The `iptables` interface is one of several kernel modules based on the netfilter framework; others include a module that handles Network Address Translation (which enables multiple machines to share one public IP address), and the module that implements and supports connection tracking. Throughout the rest of this chapter, I will collectively refer to this as `iptables`, because that is the interface that is most commonly associated with modern Linux firewalls and packet filtering.

The `iptables` interface and the netfilter framework are actually the fourth generation of Linux packet filtering solutions. The original Linux packet filtering implementation, `ipfw`, was liberated from BSD-based systems and was introduced in Linux by Alan Cox in the Linux 1.1 kernel, and was designed to support the creation of simple IP firewalls and routers through packet inspection and filtering. The `iwfwadmin` tool and associated `ipfw` changes, which simplified creating `ipfw`-based firewalls, was added to the Linux 2.0 kernel and makes up the second generation. The third generation of Linux packet filtering, consisting of a major rewrite of the entire Linux networking layer and introducing the user-space `ipchains` tool, was introduced in the 2.1 kernel series. The current netfilter framework and `iptables` interface were introduced in the 2.4 kernel, and have been the standard mechanism for packet filtering, network address and port translation, and general packet manipulation (often referred to as packet mangling) in the 2.6 series of Linux kernels.

Linux packet filtering works by inspecting incoming and outgoing packets and acting upon them based on filtering rules that have been loaded into the netfilter framework's filter table by the iptables command. By default, the iptables command supports three default sets of rules, known as *chains*, for filtering network packets using the information stored in the `iptables` filter table. These default chains are the INPUT, OUTPUT, and FORWARD chains. The rules in the INPUT chain are used to examine and process incoming packets intended for ports on the local machine. The rules in the OUTPUT chain are for examining and processing outgoing packets that are being sent from the local machine. The rules in the FORWARD chain are used to examine and process packets that are being routed through the local machine.

Each of the default filtering rule chains can have its own set of filtering rules. You can also define other sets of rules and use them for your own purposes. Many modern Linux and other UNIX-like systems come with predefined INPUT, OUTPUT, and FORWARD rule chains and automatically load them at boot time. As discussed later in this chapter, a variety of graphical and command-line software is available for all Linux distributions to make it easy to define your own packet filtering rules.

Other netfilter-based modules use packet-matching tables other than the filter table. The NAT module uses the NAT table, which contains three built-in rule chains: OUTPUT, POSTROUTING, and PREROUTING. Specialized packet manipulation operations use the mangle table, which contains prebuilt FORWARD, INPUT, OUTPUT, PREROUTING, and POSTROUTING chains. The connection tracking module uses the raw table, which contains preconfigured OUTPUT and PREROUTING chains.

You must have superuser privileges to examine, create, or modify any netfilter-based rule chains. You can do this by putting `iptables` commands in a script that is executed as part of the system's boot process or by using a command such as sudo as a normal user to run the `iptables` commands with root privileges.

Installing and Configuring a Firewall

As mentioned in the previous section, many different software packages are available to help you configure and activate a firewall on your Ubuntu or Kubuntu system. These packages include the default package that is installed on Ubuntu and Kubuntu systems, `ufw` (the Uncomplicated Firewall), Lokkit, Firestarter, Fwbuilder, Guarddog, and many more.

The default firewall package for Ubuntu and Kubuntu systems, `ufw`, is relatively easy to configure from the command line, but is even easier to configure using its associated graphical configuration utility, `gufw`. Unfortunately, this utility is not installed by default on Ubuntu or Kubuntu systems. You will have to install it using your favorite package management tool, as discussed in Chapter 23. This graphical configuration tool is GNOME-oriented but can be installed on a Kubuntu system if you don't mind also installing a number of GNOME-related packages that it requires. Personally, I think it's worth it, and therefore the rest of this section focuses on using the graphical `gufw` utility.

Installing the `gufw` graphical configuration utility adds the Firewall Configuration menu entry to the System ➪ Administration menu on a Ubuntu system, the K Menu ➪ System menu on a Kubuntu 8.04 system, or to the K Menu ➪ Applications ➪ System menu on a Kubuntu 8.01 or better system. Select this menu entry to begin configuring your firewall. After entering your password, the initial `gufw` dialog appears, as shown in Figure 25-26.

The first thing that you will want to do when configuring a firewall is to enable that firewall on your system. Click the Firewall enabled checkbox to do so. The shield graphic shown in Figure 5-26 will be displayed in green, helping you identify the fact that the firewall is now enabled and that your system is secure if you use the current configuration. The other configuration options on the gufw dialog will also be enabled, as shown in Figure 5-27.

The `gufw` tool provides several tabs that make it easy for you to create various types of rules without having to understand the syntax of the kernels `iptables` rules. Once you enable a firewall, all incoming traffic is disabled by default, which is secure but probably not what you want if you plan to support network services such as incoming SSH connections, incoming FTP, and so on. Note that you do not need to select the Allow incoming radio button shown in Figure 25-27 to enable incoming traffic. The radio buttons at the top left of the `gufw` dialog identify the default behavior of your system. Any rules that you subsequently define represent exceptions to that rule. For example, to quickly create a rule that allows incoming SSH connections, make sure that the Simple tab is selected in the Add a new rule section, enter ssh in the text entry field for this rule, and make sure that you are allowing traffic in both directions, as shown in Figure 25-28.

FIGURE 25-26

The initial gufw dialog



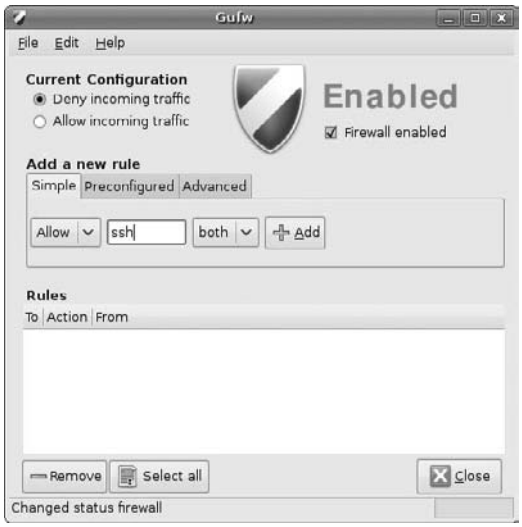
FIGURE 25-27

Enabling your firewall



FIGURE 25-28

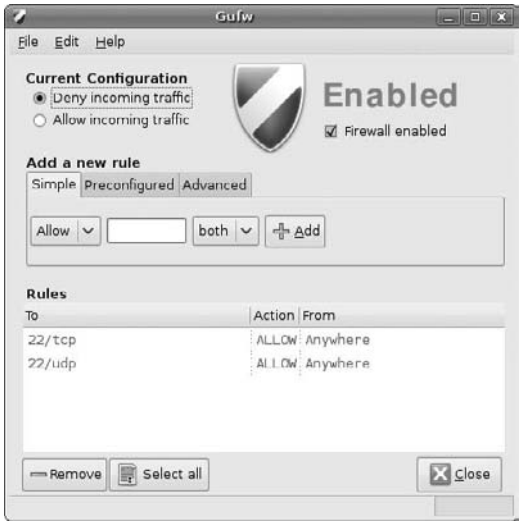
Allowing inbound/outbound SSH connections through your firewall



Click Add to add the new rule. The rule is added to the list of rules shown in the bottom of the gufw dialog, as shown in Figure 25-29, and your firewall will now allow both incoming and outbound SSH traffic.

FIGURE 25-29

Rules for allowing inbound/outbound SSH connections



The rules shown at the bottom of the `gufw` dialog refer to the ports associated with various services, which is not very user friendly. If you aren't sure about the service that a rule is associated with, you can use the `grep` command to look up the port that it uses in the file `/etc/services`. For example, the following output shows how to look for the services associated with port 22 in this file:

```
$ grep -w 22 /etc/services
ssh             22/tcp          # SSH Remote Login Protocol
ssh             22/udp
```

The other rules tabs in the `gufw` utility make it easy to define other types of rules. The Preconfigured tab provides a drop-down menu number of default rules that you can use to enable (or disable) various services, as shown in Figure 25-30.

FIGURE 25-30

Preconfigured rules for various services



You can also use the Preconfigured tab to define other types of rules by selecting them from the type drop-down menu at the left. For example, Figure 25-31 shows a rule that enables NFS traffic but also limits the bandwidth used by that service.

If you need more complex rules than the ones that can be created from the Simple or Prefined tabs, the Advanced tab, shown in Figure 25-32, provides a general mechanism for defining more sophisticated rules.

The Advanced rules tab makes it easy for you to add host or network-specific rules that enable traffic from specific hosts or networks to specific ports on your system, and vice versa.

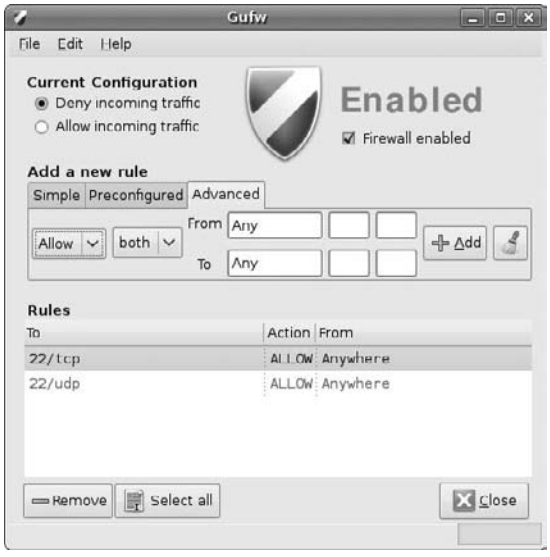
FIGURE 25-31

Defining advanced rules



FIGURE 25-32

Defining advanced rules



Once you have finished defining the rules for the traffic that your firewall will allow or deny, click Close to save your rules and terminate the `gufw` utility. If you're curious, you can use the `sudo iptables -L` command to see what the rules that you have just created look like, which make you eternally grateful for a tool like `gufw`.

Summary

This chapter explained the basics of setting up wired and wireless network connections on your Ubuntu or Kubuntu system. Both Ubuntu and Kubuntu systems provide easy-to-use graphical applets that simplify configuring these types of interfaces, eliminating the traditional command-line magic required when you are using the `ifconfig` command. This chapter also provided a basic discussion of securing your Ubuntu system and related networking hardware, and also discussed setting up a simple firewall on your Ubuntu system.

Chapter 26 discusses using wireless technologies on your Ubuntu system in more detail, providing a detailed explanation of setting up and using `Wicd`, an alternative to the Network Manager software that was discussed in this section. Chapter 26 also includes a detailed explanation of how to use Windows drivers on your Ubuntu or Kubuntu system to make new or problematic Ethernet and other devices work with your new Linux system.

Chapter 26

Going Wireless

Today's mandatory connectivity requirements would be a farce without the ubiquity of wireless connections, whether these are cell phones leveraging a distant tower, your cell phone or PDA syncing to your PC via Bluetooth or infrared, or your laptop using wireless network connections via an access point hidden in the linen closet in your hotel. It is incredibly liberating to be able to travel without carrying every cable and adapter that could possibly connect device A to device B, especially when airport security seems to frown on suitcases full of cables, power supplies, and connectivity hardware last seen in the laboratory of Victor von Frankenstein (no relation to the author).

Like many other high-tech domains, wireless networking has its own vocabulary and a bewildering set of alphabetically similar but conceptually unique acronyms, protocols, and so on. The first section of this chapter provides a basic overview of wireless networking and associated terminology today. Subsequent sections of this chapter explain how to configure wireless devices that are directly supported by Linux, explain how to work with wireless devices that are not supported by Linux (an impressive trick), and provide tips and tricks for using Microsoft Windows drivers with your Ubuntu or Kubuntu Linux system.

IN THIS CHAPTER

Understanding wireless networking terminology

Installing and using Wicd

Using Windows drivers

Troubleshooting Windows drivers

Overview of Wireless Technologies

Wireless networking is amazingly useful, both on the job (so that you never have to play rock/paper/scissors over a limited number of Ethernet cables in a conference room), and even more so when traveling, enabling you to connect laptops, phones, and PDAs to the Internet in libraries, coffee shops, hotels, and many other public places. Wireless networking also provides an excellent mechanism for connecting networks of computers that are located in inaccessible locations or in existing structures where running new or additional cabling is a problem. Modulo security concerns, wireless networks can prevent the need to drill additional holes through Mount Vernon, the Vatican, or the Louvre.

As used in this chapter, wireless networking refers to technology that enables a computer to communicate using standard network protocols, but without network cabling. Most commonly, wireless networks are local area networks (LANs) whose members use industry standards such as Institute of Electrical and Electronics Engineers (IEEE) 802.11 to communicate. The original 802.11 standard was released in 1997 and defines the frequencies, data rates, and media access methods used to communicate between the participants in a wireless network. Subsequent standards have defined a potentially bewildering variety of 802.11X standards, which I'll discuss a bit later in this section.

Regardless of the frequency and speed at which participants communicate, there are two basic types of wireless networks: ad hoc wireless networks and managed wireless networks. The characteristics of each (and their differences) are:

- **Ad hoc networks:** Also referred to as peer-to-peer wireless networks, consist of some number of computers that each have a wireless networking interface card and that communicate directly with all of the other wireless computers on that network. This enables them to share files and printers, but will not provide access to wired or Internet network resources unless one (or more) of the computers is also connected to those other network resources and is configured to serve as a gateway, bridge, or router to the other network(s). Acronym fans may see ad hoc networks referred to as IBSS (Independent Basic Service Sets) during the computer bowl.
- **Managed networks:** Also referred to as infrastructure wireless networks, these use an access point (sometimes also referred to as a base station) to manage communication between multiple wireless computers, acting much like a hub or switch for the wireless network. An access point is also typically connected to a wired network such as the Internet, and serves as the bridge or gateway between the wireless and wired networks. Access points are traditionally dedicated hardware devices (sometimes referred to as a HAP, a hardware access point), such as Apple's Airports or devices from companies such as DLink, LinkSys, NetGear, Cisco, and others, but can also be software solutions (referred to as SAPs, software access points) that run on a computer that is equipped with both wired and wireless network connections. Acronym fans may see managed networks referred to as a BSS (Basic Service Set; *N* computers and one access point) or ESS (Extended Service Set; *N* access points and *N* computers forming one subnet) during the computer bowl.

Large areas that require wireless coverage can accomplish this by providing multiple access points and supporting roaming, which is the ability of a user's connection to transfer from one access point to another. This is usually invisible to the user, though some access points require passwords or other authentication when moving between access points. Many networking hardware vendors also provide specialized hardware known as *extension points* that amplify the signal and therefore extend the range of an existing access point. Whether or not hardware access points support roaming or extension points is dependent on the hardware manufacturer. As with most networking hardware, a good rule of thumb is that buying all of your network hardware from the same vendor increases your chances for compatibility.

Agreement between the low-level wireless networking standards that the computers in your wireless network use is the most important aspect of setting up a successful wireless network. If the computers and access points can't communicate in the first place, you aren't going to be setting up much of anything. Though IEEE 802.11 is a published standard for wireless communication, several different 802.11 protocols and associated standards exist, not all of which are compatible with each other. The following is a current list of popular 802.11 standards, communication speeds, and compatibility promises:

- **802.11a:** A standard for 802.11 communications using the regulated 5.0 GHz frequency, offering maximum communication speeds of 25 to 54 megabits per second. It is not compatible with any other 802.11 standard.

- **802.11b:** A standard for 802.11 communications using the unregulated 2.4 GHz frequency, offering maximum communication speeds of 5 to 11 megabits per second. The 802.11b standard is forward compatible with the 802.11g and proposed 802.11n standards.

TIP **Some vendors offer wireless cards known as *dual-band* cards, which support both 802.11a and 802.11b, operating in the frequency range and with communications speeds dictated by the type of connection that they establish.**

- **802.11g:** A standard for 802.11 communications using the unregulated 2.4 GHz frequency, offering maximum communication speeds of 25 to 54 megabits per second. The 802.11g standard is backward compatible with the 802.11b standard and forward compatible with the proposed 802.11n standard.
- **802.11n:** A standard for 802.11 communications using the unregulated 2.4 GHz frequency, offering maximum communications speeds of 100 to 300 megabits per second. The 802.11n standard is designed to be backward compatible with the 802.11b and 802.11g standards.

CAUTION **At the time that this book was written, the *n* in 802.11n should stand for "not quite a standard." The 802.11n standard is still under discussion and is not slated for official publication until early 2009, depending on whom you ask. However, the lack of an official, frozen standard hasn't stopped vendors from releasing supposed 802.11n hardware—you will see many wireless cards advertised as 802.11 Draft-N cards to recognize the fact that they adhere to an early version of the standard. Until the official 802.11n standard is published, a good rule of thumb for compatibility between supposed 802.11n devices is "buy everything from the same vendor." Most vendors' hardware will, after all, work with their other hardware. Worst case, you only have one vendor to complain to.**

Many wireless cards are advertised as 802.11 b/g or 802.11 b/g/n cards, because they can communicate with all of these networks. When using cards that support all of these standards, the speed of your connection is determined by the speed/type of the access point that they establish a connection with.

Wireless networks also support a variety of security solutions, ranging from network names to encryption mechanisms. The most common of these are the following:

- **MAC (Media Access Control) filtering:** A security mechanism that requires that access points be programmed with lists of the systems that can connect to them, identified by hardware Ethernet address. MAC security only prevents against unauthorized connections to an access point—it does not secure those communications, once established.
- **SSID (Service Set Identifier):** An SSID acts as a simple password by providing a unique identifier for a specific wireless network. Access points with a specific SSID can be configured to disallow access to anyone who does not provide that SSID when negotiating the initial connection. SSID security only prevents unauthorized connections to an access point—it does not secure those communications, once established. An Extended Service Set Identifier (ESSID) is just an SSID that is (or can be) used on multiple access points to identify the same network.
- **WEP (Wired Equivalent Privacy):** A security mechanism that requires supplying a 48-, 64-, or 128-bit security key when negotiating a connection to an access point. This key is used for encrypting and decrypting wireless communications. If this key is the same as that used by the access point, the two can communicate successfully. WEP security protects against unauthorized access and also provides secure wireless communication, because all communication packets are encrypted. Although still quite common, WEP (Wired Equivalent Privacy) keys can be cracked in a few minutes using popular open source packages such as Aircrack-ng (`www.aircrack-ng.org/doku.php`) and Kismet (`www.kismetwireless.net/index.shtml`).

- **WPA (Wi-Fi Protected Access):** A security mechanism that uses a Temporary Key Integrity Protocol to replace WEP and provides enhanced security on existing hardware. WPA uses a key server or pregenerated key set to encrypt communications on a per-packet basis. Multiple WPA standards exist: WPA, which was developed using a draft of the IEEE 802.11i security standard, WPA2, which was developed using an approved version of the 802.11i security standard, WPA Personal, which uses preshared key mode (PSK) and stores keys in access points, and WPA Enterprise, which uses a networked authentication server and supports a number of EAP (Extensible Authentication Protocol) security mechanisms.

All of these security measures can be used together, in any combination, or separately. Only WEP and WPA provide security for the packets that are being transmitted, but typing in a 128-bit key each time you configure a new wireless interface is both error-prone and incredibly tedious. Many sites, therefore, use MAC and SSID/ESSID security to establish connections, and then protect transmitted packets by using standard SSH tunneling or VPN technologies on top of the wireless communication layer.

The Linux kernel includes drivers for many popular PCI wireless cards, PCMCIA wireless adapters, and USB wireless adapters. On Ubuntu systems, wireless Ethernet interfaces are configured in the same way as wired Ethernet interfaces, which I'll explain in the next section.

If no Linux driver is available for your wireless card, all is not lost. Thanks to the cleverness of the folks in the open source community, a Linux kernel module and related utilities, called `ndiswrapper`, enables you to use the Microsoft Windows drivers for much of the networking hardware that is not yet natively supported by Linux. Installing these utilities, installing Windows drivers under Linux, and configuring your system to automatically start loading the kernel module is discussed later in this section.

Using systems that have both wired and wireless interfaces is quite common nowadays. Even though wireless specifications such as 802.11g and the upcoming 802.11n provide high communication speeds, they can't match the speed of the high-speed wired connections such as Gigabit Ethernet connections. It's, therefore, a shame not to use your system's wired connection whenever possible. However, it's tough to travel and get work done while dragging along hundreds of miles of Ethernet cable—the convenience and flexibility of wireless network access far outweighs the slower performance that it provides. The Network Manager and the Ubuntu and Kubuntu applets that simplify configuring it make it relatively easy to use a wireless network interface from a variety of locations, and were discussed in detail in Chapter 25. However, we're talking about Linux here, so there are always alternatives. One extremely impressive alternative to the Network Manager is the `wicd` package, which is discussed in the next section.

Using Wicd

Wicd is an alternate network manager for Linux systems that makes it easy to establish connections to wired and wireless networks. Wicd is often easier to use and is frequently more successful at establishing connections to secure networks than the traditional Linux Network Manager. The Network Manager is a fine piece of software that usually works—however, when it doesn't, it is a total PITA (and I'm not referring to the popular Middle Eastern bread). As a replacement for the Network Manager and related applets, Wicd conflicts with the Network Manager, and only one of these network connection management packages can be installed on a system at a time. Luckily, as discussed in the next section, installing Wicd on

a system where the Network Manager is installed will automatically purge the Network Manager for you. No fuss, no muss.

Wicd is a GTK application, unlike the graphical interfaces for the Network Manager that are provided by applets which are specific to the desktop environments used on both standard Ubuntu (GNOME) and Kubuntu (KDE 3 and KDE 4) systems, which were discussed in Chapter 25. As a GTK application, Wicd can therefore be used on both GNOME and KDE-based systems, and only drags a minimal amount of package baggage along with it when installed on Kubuntu systems.

The next sections explain how to install, use, and configure Wicd to meet your needs. It's really a great application, which I use on all of my laptop systems (and any desktops with both wired and wireless network interfaces) rather than the Network Manager.

Installing Wicd

Both Ubuntu and Kubuntu systems are currently committed to the Network Manager and related applets, so in order to install Wicd, you will first have to modify the list of sources that provide repositories from which your Ubuntu or Kubuntu systems can install software. You can do this using Ubuntu's System ➪ Administration ➪ Software Sources application, using the Sources icon in Kubuntu's Adept Manager, or by directly modifying your `/etc/apt/sources.list` file. You will need to add the following entry in any of these locations:

```
deb http://apt.wicd.net intrepid extras
```

If you are using a version of Ubuntu or Kubuntu other than intrepid (8.10), you should replace intrepid with the name of the version of Ubuntu or Kubuntu that you are using (jackalope, hardy, etc.). If you are using Synaptic or Adept, those package management applications will automatically update the lists of available packages that they know about. If you directly modified your `/etc/apt/sources.list` file, you will need to tell whatever package management software you are using to update its package lists. If you are using aptitude to install software, you can do this by executing the command `sudo aptitude update` from any command-line environment. If you have problems updating your package management application's list of available packages, see Chapter 23 for more information about updating package lists in whatever package manager you are using.

Before installing Wicd, you should also import its privacy key so that your package manager won't warn you about the dangers of installing software from unauthorized sources. You can import this key from any command-line environment by executing the following command:

```
wget -q http://apt.wicd.net/wicd.gpg -O- | sudo apt-key add -
```

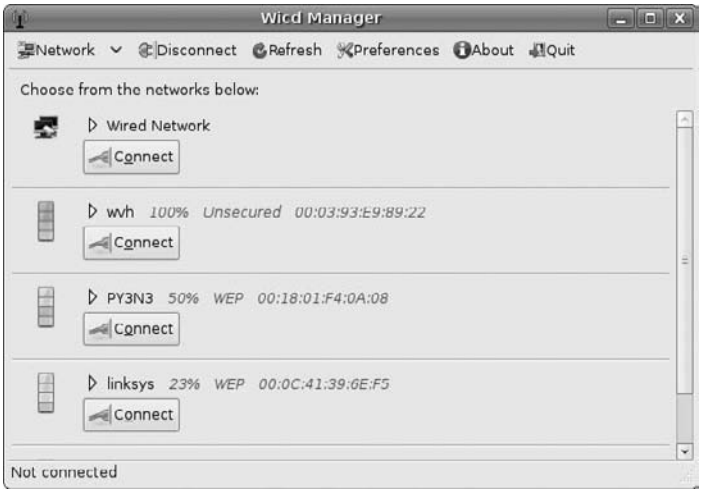
After importing the Wicd key, you can use your favorite package management tool to install the `wicd` package. As a replacement for the Network Manager and its applet buddies, Wicd conflicts with the Network Manager, and only one of these network connection management packages can be installed on a system at a time. However, as mentioned previously, installing Wicd on a system where the Network Manager is installed will automatically purge the Network Manager and related applets for you. Installing Wicd will install both the Wicd daemon (`/usr/bin/wicd`) and a Python-based client application for your Ubuntu system's Notification Area or your Kubuntu system's System Tray. This client application will start automatically the next time that you log in on your system, but if it doesn't start automatically when you install Wicd, you can execute it manually by executing the command `/usr/bin/wicd-client&` from any command-line environment.

Using Wicd

To start the client portion of Wicd, click on its tray icon to display a screen like the one shown in Figure 26-1.

FIGURE 26-1

The Wicd Manager dialog on Ubuntu



This initial dialog should show an entry for the wired connection on your system and any wireless networks that have been detected. The wireless networks that are listed depend on those that are detected in your location—if the screen that you see when you first start Wicd looks exactly like the one shown in Figure 26-1, then I'm flattered (and a little creeped out) that you're reading this book from somewhere in my house.

To demonstrate that Wicd does indeed work fine on Kubuntu systems, Figure 26-2 shows the Wicd Manager running on a Kubuntu laptop. (The rest of the screenshots in this section were taken on a Ubuntu laptop, because the interfaces on Ubuntu and Kubuntu really are identical aside from the window decorations and styles.)

As you can see from these figures, Wicd initially identifies all available wired and wireless connections, but it does not automatically connect to any of them. To connect to any of these, click the Connect button below that entry. If you attempt to connect to a network that requires authentication, you will immediately see a dialog like the one shown in Figure 26-3, which tells you that the connection requires encryption.

If you see the dialog shown in Figure 26-3, click OK to dismiss it and click the right arrow beside the name of the network that you were attempting to connect to. This expands the entry for that network, exposing the Advanced Settings button that you must use to define the encryption/authentication information that is required to establish a connection to that network, as shown in Figure 26-4.

FIGURE 26-2

The Wicd Manager dialog on Kubuntu

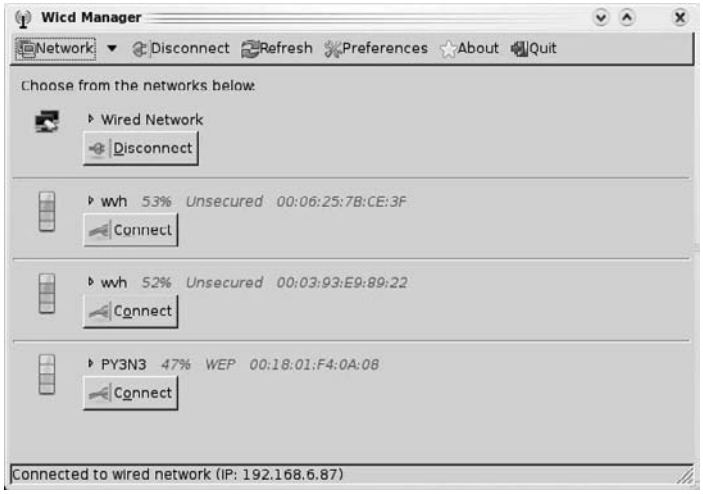


FIGURE 26-3

Attempting a connection that requires a password or key

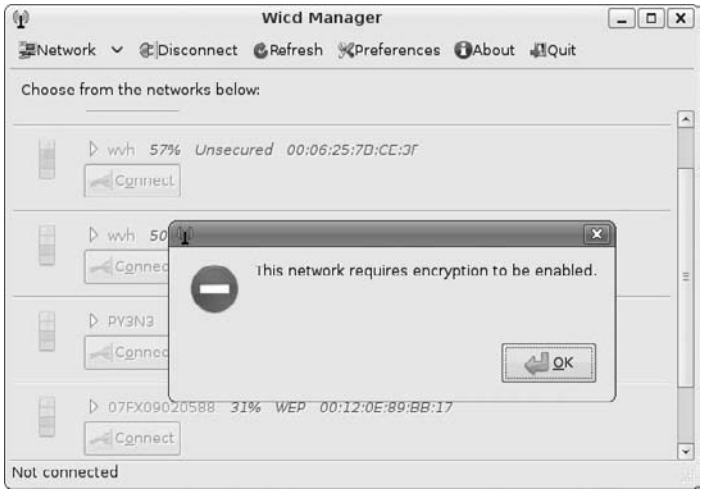
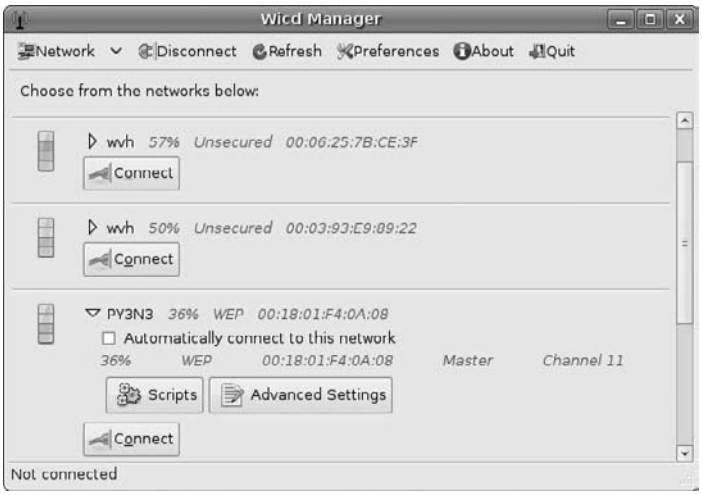


FIGURE 26-4

Exposing configuration options for a given wireless network



Click Advanced Settings to display the dialog that enables you to define the authentication or encryption information that you must supply in order to access a specific network. A dialog like the one shown in Figure 26-5 appears.

FIGURE 26-5

Configuring encryption/authentication for a specific network connection



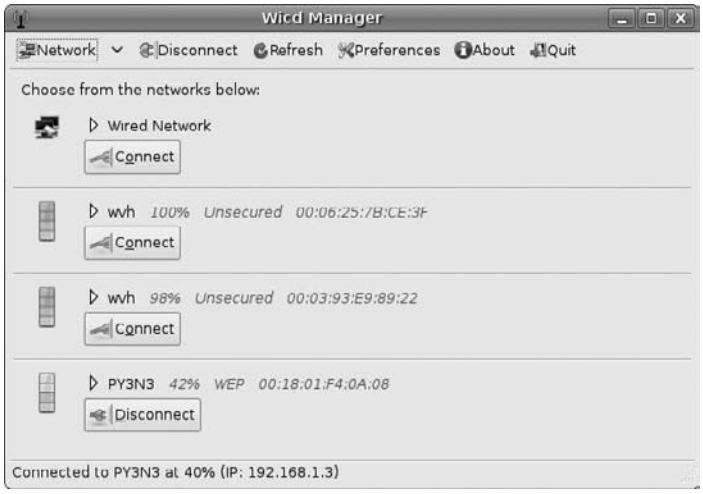
Make sure that the Use Encryption checkbox is selected, and click the drop-down list near the bottom of this dialog to display your encryption and authentication options. Select the mechanism that is used by the wireless network that you want to connect to, and enter the information that it requires in the fields that are displayed. For example, Figure 26-5 shows the area for entering a 40/128-bit WEP key for a network that uses WEP.

TIP **Each wireless network shown in the dialog in Figure 26-1 or 26-2 identifies the type of encryption or authentication that it uses in the heading for that wireless network. For example, in Figures 26-1 and 26-2, you can see that the wireless network `wvh` does not use authentication or encryption, while the wireless network PY3N3 uses WEP.**

After entering the information required by the encryption or authentication mechanism associated with the wireless network that you want to connect to, click OK to close the Advanced Setting dialog and return to a screen like the one shown in Figure 26-4. If you want to try to connect to this wireless network each time you start your system, select the Automatically connect to this network option. Click Connect to connect to the specified wireless network. You will see status messages and a progress bar at the bottom of the Wicd Manager dialog, and the Wicd Manager will display a screen like the one shown in Figure 26-6 once the connection has been established successfully.

FIGURE 26-6

A successful connection to a wireless network



This dialog displays the IP address and transmission rate for your new connection at the bottom of the Wicd Manager window. You can now close the Wicd Manager dialog by clicking the close box in its upper right-hand corner.

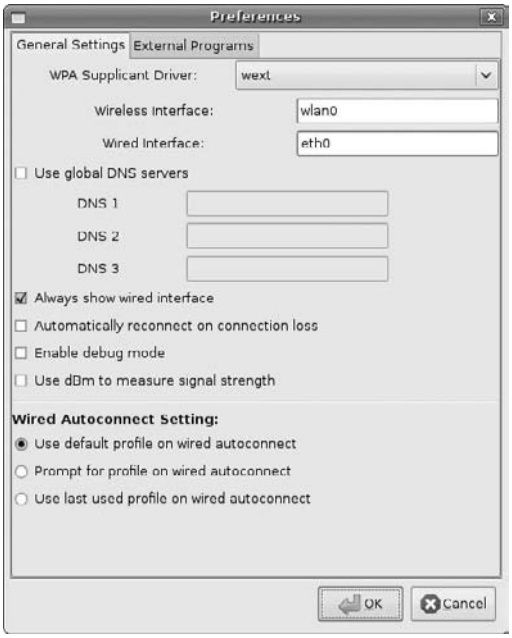
That's really all there is to it, except for some optional Wicd configuration that you may or may not want to do. Configuring Wicd is discussed in the next section.

Configuring Wicd

To configure Wicd, select the Preferences menu option from the Wicd Manager to display its configuration dialog, shown in Figure 26-7.

FIGURE 26-7

Setting preferences for Wicd



In most cases, you won't need to modify Wicd's basic settings, but this dialog provides some useful options. The handiest of these is the ability to identify the names of the Ethernet devices associated with wired and wireless devices on your system. If you insert or change wireless adaptors while Wicd is running, you will probably need to modify the name of the wireless Ethernet device. After inserting a new wireless interface, you can determine the name of that interface by using the `iwlist` command, which is discussed in the next section. Sample output from this command looks like the following:

```
$ iwconfig
lo        no wireless extensions.

eth0      no wireless extensions.

irda0     no wireless extensions.

pan0      no wireless extensions.

wlan0     IEEE 802.11b+/g+  ESSID:"wvh"  Nickname:"acx v0.3.36"
```

```
Mode:Managed  Frequency:2.437 GHz  Access Point: 00:06:25:7B:CE:3F
Bit Rate:5.5 Mb/s   Tx-Power=15 dBm   Sensitivity=1/3
Retry min limit:7   RTS thr:off
Power Management:off
Link Quality=31/100  Signal level=3/100  Noise level=0/100
Rx invalid nwid:0  Rx invalid crypt:0  Rx invalid frag:0
Tx excessive retries:0  Invalid misc:0   Missed beacon:0
```

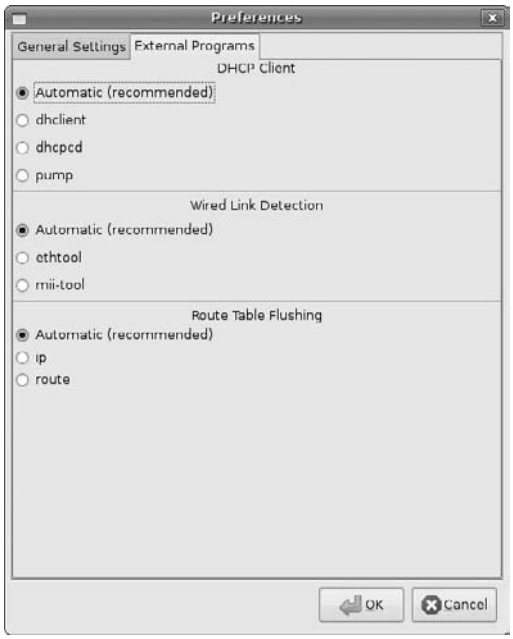
In this case, the new wireless device on this system is `wlan0`, which you should then enter as the name of your wireless device in the dialog in Figure 26-7. If you change the name of your wireless (or wired) device, you can then click close to close the Preferences dialog, and select the Refresh menu option in the Wicd Manager dialog to re-probe your networked environment and identify the wireless networks that can be detected using the new device.

The other option that I find useful in the dialog shown in Figure 26-7 is the Always show wired interface option, which should be selected to always list your system's wired interface in the main Wicd manager dialog.

To see additional Wicd settings, click the External programs tab to display the Preferences pane shown in Figure 26-8. Although you will rarely have to modify any of the settings on this dialog, doing so enables you to identify a specific application that you want to use as a DHCP client, to detect wireless links, and to purge the routing table when switching network connections.

FIGURE 26-8

Setting application preferences for Wicd



After making your changes or simply examining Wicd's current configuration, click OK to close the preferences dialog and return to the Wicd Manager.

Command-Line Tools for Wireless Networking

The `wireless-tools` package, which is installed on Ubuntu and Kubuntu systems by default, provides several commands for probing and configuring wireless interfaces and networks. You can execute the `man wireless` command to get information about the `wireless-tools` package, learn how and where it is installed by various Linux distributions, and view a list of the commands that this package provides.

The `wireless-tools` commands that I find most valuable are the `iwconfig` and `iwlist` commands. The `iwconfig` command displays information much like the `ifconfig` command, providing information about the wireless status of all or specific Ethernet interfaces on a system. The following example shows the output of the `iwconfig` command, querying the wireless Ethernet interface that was installed in the previous section:

```
$ iwconfig wlan0
wlan0     IEEE 802.11b+/g+  ESSID:"wvh"  Nickname:"acx v0.3.36"
          Mode:Managed  Frequency:2.437 GHz  Access Point: 00:06:25:7B:CE:3F
          Bit Rate:5.5 Mb/s   Tx-Power=15 dBm   Sensitivity=1/3
          Retry min limit:7   RTS thr:off
          Power Management:off
          Link Quality=31/100  Signal level=3/100  Noise level=0/100
          Rx invalid nwid:0  Rx invalid crypt:0  Rx invalid frag:0
          Tx excessive retries:0  Invalid misc:0   Missed beacon:0
```

As you can see, this provides much of the same information as that provided by the `ifconfig` command, but it also includes wireless-specific information. The `iwconfig` command can also be used to set various wireless configuration options from the command line—see the online man page for `iwconfig` (using the command `man iwconfig`) for more information.

The `iwlist` command enables you to scan for wireless networks and examine various characteristics of a wireless interface. For example, scanning for wireless networks around my home using the `iwlist wlan0 scan` command displays the following information:

```
$ sudo iwlist wlan0 scanning
wlan0     Scan completed :
          Cell 01 - Address: 00:18:01:F4:0A:08
                    ESSID:"PY3N3"
                    Mode:Master
                    Frequency:2.462 GHz (Channel 11)
                    Quality=37/100  Signal level=12/100  Noise level=0/100
                    Encryption key:on
                    Bit Rates:1 Mb/s; 2 Mb/s; 5.5 Mb/s; 6 Mb/s; 9 Mb/s
                              11 Mb/s; 12 Mb/s; 18 Mb/s; 24 Mb/s; 36 Mb/s
                              48 Mb/s; 54 Mb/s
          Cell 02 - Address: 00:03:93:E9:89:22
                    ESSID:"wvh"
                    Mode:Master
                    Frequency:2.427 GHz (Channel 4)
                    Quality=52/100  Signal level=34/100  Noise level=0/100
                    Encryption key:off
```

```
                    Bit Rates:1 Mb/s; 2 Mb/s; 5.5 Mb/s; 6 Mb/s; 9 Mb/s
                              11 Mb/s; 12 Mb/s; 18 Mb/s; 24 Mb/s; 36 Mb/s
                              48 Mb/s; 54 Mb/s
          Cell 03 - Address: 00:06:25:7B:CE:3F
                    ESSID:"wvh"
                    Mode:Master
                    Frequency:2.437 GHz (Channel 6)
                    Quality=57/100  Signal level=40/100  Noise level=0/100
                    Encryption key:off
                    Bit Rates:1 Mb/s; 2 Mb/s; 5.5 Mb/s; 11 Mb/s
```

This command is useful for verifying that a wireless interface is actually working, and that it can communicate with an access point. The information produced by the `iwlist` command is very useful in public areas, where you will usually detect more wireless networks, as in the following example from one of my laptops while I was on vacation:

```
$ sudo iwlist eth0 scan
Password:
eth0      Scan completed :
          Cell 01 - Address: 00:09:5B:36:22:38
                    ESSID:"Wireless"
                    Mode:Master
                    Frequency:2.412 GHz (Channel 1)
                    Signal level:20/153  Noise level:12/153
                    Encryption key:off
                    Bit Rates:1 Mb/s; 2 Mb/s; 5.5 Mb/s; 11 Mb/s
          Cell 02 - Address: 00:12:17:A6:E8:CC
                    ESSID:"linksys"
                    Mode:Master
                    Frequency:2.437 GHz (Channel 6)
                    Signal level:83/153  Noise level:13/153
                    Encryption key:off
                    Bit Rates:1 Mb/s; 2 Mb/s; 5.5 Mb/s; 11 Mb/s
          Cell 03 - Address: 00:12:17:A6:E8:CC
                    ESSID:"linksys"
                    Mode:Master
                    Frequency:2.437 GHz (Channel 6)
                    Signal level:81/153  Noise level:13/153
                    Encryption key:off
                    Bit Rates:1 Mb/s; 2 Mb/s; 5.5 Mb/s; 11 Mb/s
          Cell 04 - Address: 00:12:17:79:CF:D3
                    ESSID:"linksysR79cfd3"
                    Mode:Master
                    Frequency:2.437 GHz (Channel 6)
                    Signal level:49/153  Noise level:13/153
                    Encryption key:off
                    Bit Rates:1 Mb/s; 2 Mb/s; 5.5 Mb/s; 11 Mb/s; 18 Mb/s
                              24 Mb/s; 36 Mb/s; 54 Mb/s
```

In this case, the information provided by the `iwlist` command was very useful in identifying available wireless access points, their names, and whether they were configured for authentication/encryption. The `iwlist` command can also query specific wireless configuration details—see the online man page for `iwlist` (using the command `man iwlist`) for more information.

Installing and Using Windows Networking Drivers

Device driver support is one of the few areas in which Linux lags behind Microsoft Windows. I hope that my making this statement doesn't induce Linus Torvalds to give me a call and demand the return of my official Linux Bigot decoder ring, but it's a fact. The reason that there are not Linux drivers for every piece of random hardware is not because of any problem or limitation in Linux but rather because device manufacturers target their largest potential markets when writing device drivers for their hardware, and at the moment, the main market for random PC hardware is Microsoft Windows systems.

Clever folks in the open source community are writing new device drivers every day, and each new Linux kernel supports more hardware than ever before. However great that is, it doesn't help much when you need to use a specific Windows device today. Luckily, a Linux kernel module and related utilities, called `ndiswrapper`, enables you to use the Microsoft Windows drivers for much of the networking hardware that is not yet natively supported by Linux. How cool is that? Thanks to the `ndiswrapper` package, you can use Windows' own drivers on your superior Ubuntu Linux systems until a native Linux driver becomes available.

One obvious question is, "Why not just buy hardware that works with Linux in the first place?" To be honest, that's my general approach, but there are some very good reasons for wanting (or needing) to use hardware that is supposedly Windows-only. The most significant of these are the following:

- You already have a wireless network adapter that is not yet supported by Linux. Why buy another adapter if you don't have to?
- Dual-boot laptops may include built-in wireless hardware that is not supported by Linux. Similarly, if you are given a laptop as part of your job along with PCMCIA or USB network adapters that are not supported by Linux, using the Windows drivers with Ubuntu enables you to use a real operating system for work without requiring that you spend any money.
- Some wireless protocols, most specifically 802.11a, are poorly supported under Linux but are still used by many businesses. Because 802.11a is also largely extinct, I doubt that this will change. Using the Windows drivers for your 802.11a hardware enables you to run Linux at work, using existing hardware.

The next few sections explain how to install this package and a handy graphical configuration utility, describes how to install a Windows driver on your Ubuntu system, and provides some tips and tricks for debugging problems and resolving problems using Windows drivers on your Ubuntu system.

TIP **Though primarily designed for use with Microsoft Windows drivers for network cards, the `ndiswrapper` package has been used with Windows drivers for many other devices, such as USB serial port devices, wired Ethernet cards, and so on. (For a complete list of devices that people have gotten to work with Windows drivers and the `ndiswrapper` package, see `http://ndiswrapper.sourceforge.net/mediawiki/index.php/List`). If you have a device that came with Windows drivers that you can't get to work natively under Linux, the `ndiswrapper` package is certainly worth a try and is only a package management tool session install away.**

Installing NDIS Wrapper and Friends

The `ndiswrapper` Linux kernel module is provided with every Ubuntu and Kubuntu Linux kernel, but you'll still want to install some related utilities to simplify installing and configuring Microsoft Windows drivers. The `ndisgtk` package provides a graphical utility that simplifies installing and configuring

Windows drivers; the `ndiswrapper-utils` package provides the underlying utilities that it uses, as well as a command-line utility that you can use to install Windows drivers if you're not a GUI fan. These aren't installed by default on Ubuntu or Kubuntu systems because not everyone needs them, and why waste disk space if you don't have to?

See Chapter 23, "Adding, Removing, and Updating Software," for more detailed information on installing applications using a variety of package management tools.

Installing Microsoft Windows Drivers

This section explains how to install Microsoft Windows drivers using the graphical `ndisgtk` utility, which (if you'll pardon the expression) puts a friendly wrapper around the `ndiswrapper` utilities, kernel module, and Ubuntu network configuration. You can also install Windows drivers for use with the ndiswrapper kernel module by using the command `sudo ndiswrapper -i` *`name-of-INF-file`* from any command-line environment, and then manually modifying some configuration files, but I prefer the graphical one-stop-shopping approach provided by the `ndisgtk` utility.

NOTE **Before installing Windows drivers, you must make sure that they are available on your Ubuntu or Kubuntu system. This usually involves either putting the driver CD in your system so that it is available as `/media/cdrom` (or something similar, depending on the drives in your Ubuntu system), or copying a zip archive containing the drivers to your Ubuntu system and using the `unzip` command to extract the contents of that archive file into an Ubuntu directory.**

Installing the `ndisgtk` package creates the System ➪ Administration ➪ Windows Wireless Drivers menu item on a Ubuntu System, or the System ➪ Windows Wireless Drivers menu item on Kubuntu systems, which you can select to start the `ndisgtk` application. Figure 26-9 shows the initial driver installation dialog that is displayed when you first start this application.

FIGURE 26-9

The ndisgtk startup dialog



To install a driver, click Install New Driver. This displays the dialog that enables you to select the INF file associated with a specific Window driver, as shown in Figure 26-10. An INF file is a file with a .INF or .inf file extension that contains a description of a Windows driver, its installation options, associated files, and so on.

FIGURE 26-10

The INF file selection dialog for a Microsoft Windows driver



Click the Location button to display a standard file dialog that enables you to navigate to the directory containing the Windows driver and select the appropriate INF file. After navigating to and selecting the appropriate file, click Open to select that file and close the navigation dialog. The dialog shown in Figure 26-10 redisplays, showing the name of the selected INF file. (In this particular example, I'm installing the Windows drivers for a SMB W541C wireless adapter, although you wouldn't be able to guess that from the name of the driver or INF file.)

Click Install to install the driver and associated files for use by the `ndiswrapper` kernel module. The dialog shown in Figure 31-9 redisplays, now listing the driver that you've just installed and identifying if the hardware associated with that driver was located, as shown in Figure 31-11.

FIGURE 31-11

Loading the driver and detecting associated hardware



At this point, you can click Close to terminate the Windows driver installation process. Your network connection management software, whether Wicd or a Network Manager applet, should detect the new interface and offer options for using it.

To ensure that the `ndiswrapper` kernel module loads each time you restart your Ubuntu system, the `ndiswrapper` kernel module must be listed in the file that identifies modules that should be loaded at boot time, `/etc/modules`. This should have happened when you installed the package, but you may want to check. If you need to add an entry for this module, you can do so using your favorite text editor (which you must start using `sudo` so that you can edit this file), or via a command such as the following:

```
$ sudo echo "ndiswrapper" >> /etc/modules
```

Make sure that you use two angle brackets, not just one! If you accidentally use only one angle bracket, you will overwrite the contents of your existing `/etc/modules` file with a single line containing the name of the `ndiswrapper` module, which will probably certainly cause things like your mouse and printers to misbehave.

NDIS Wrapper Tips and Tricks

As you can see from the previous section, many Microsoft Windows drivers "just work" with the `ndiswrapper` kernel module when installed correctly. Unfortunately, some don't. I won't pretend to have tested every possible piece of wireless hardware and its associated Windows drivers (though I'd be happy to try if you would just send me a sample of the hardware that I can keep), but I have used many different ones with a wide selection of wireless hardware when configuring laptops for a variety of customers, coworkers, and friends.

For me, one of the most frustrating sets of experiences has been with a variety of LinkSys WPC54G wireless PCMCIA cards, all of which I've gotten working, but only after significant driver gymnastics, head scratching, Google searches, and commiseration with fellow Linux devotees. (I was originally planning to title this section "The LinkSys WPC54G, or How I Spent My Summer Vacation," but I am trying to generalize what I've learned rather than simply provide a single set of instructions.) By the way, these are great cards that support 802.11g (and therefore 802.11b), and because I typically find LinkSys hardware to be inexpensive and reliable, I often recommend them to friends. Unfortunately, a HUGE problem with these cards is that LinkSys has used a variety of wireless chips in them over the years, so they often use different drivers, each of which seems to require its own tweaks to get it to work with the `ndiswrapper` module.

One key item is to make sure that you have the right Windows drivers in the first place. For example, some of the 54G cards use a Realtek 8180 chipset, and, therefore, only work with these drivers. Others use a Texas Instruments ACX 1*xx* chipset, while others use the Broadcom BCM43*xx*, and some newer ones are rumored to use a Marvell chipset. If you can't identify the chipset by reading the product literature, you can usually identify it by using the `lspci` utility to list all of the PCI devices on your system and by scanning the `lspci` output for anything identified as a wireless controller.

As an example of driver and chipset weirdness, the Texas Instruments ACX 1*xx* chipset is directly supported under Linux, using the `acx` and `acx_pci` kernel modules (no `ndiswrapper` module required, although I prefer the `ndiswrapper` approach), but often requires that you play around with different versions of the microcode in `/lib/firware/`*`kernel-version`*`/acx` until you find the right version and name (see the Ubuntu Forums for some suggestions). If you want to use the `ndiswrapper` version of the 54G driver with `ndiswrapper`, you have to blacklist the standard Linux kernel modules by adding their names to `/etc/modprobe.d/blacklist` so that they are not loaded during the hardware detection phase of the boot process.

If you're sure that you have the right driver, but still can't get it to work, use the Linux `dmesg` command to display recent kernel messages for error messages related to the `ndiswrapper`, the card that you're using, or any of the chipsets mentioned in the previous paragraph. This can help identify specific problems.

The WPC54G cards that use the Broadcom BCM43xx have a particularly interesting wrinkle—they require a specific microcode firmware driver that is bundled inside the Windows driver (`bcmwl5.sys`) but that must be provided as a standalone file in the `/lib/firmware` directory that the Linux kernel checks for firmware microcode. You must, therefore, use a utility called `bcm43xx-fwcutter` to extract the different firmware files from the driver, and copy them into /lib/firmware before you can use this version of the card. If you have this version of the card, you can download the latest version of this utility from the page at `http://developer.berlios.de/project/showfiles.php?group_id=4547`.

As mentioned previously, these are examples of the types of problems that you may encounter when working with a random Windows driver, random wireless hardware, and the `ndiswrapper` package. Many people seem worried about the quality of open source code, but I think that the quality of Windows drivers is just as suspect. However, as my grandfather used to say, "A dysfunctional driver in the hand is worth several driver projects that are under development in the bush." (Or something like that.)

TIP **The `ndiswrapper` developers provide as much online help as they can via the resources listed at `http://ndiswrapper.sourceforge.net/support.html`. If the people there can't solve your problem and you absolutely need to use a specific piece of "Windows-only" wireless hardware with your Linux system, a company called Linuxant (`www.linuxant.com`) provides an excellent piece of software called DriverLoader that is very similar to the `ndiswrapper` package. The Linuxant folks test their product on a huge selection of wireless hardware. They also offer Linux drivers for selected softmodem cards (modem cards that do everything in software, a traditional pain for Linux users). You can download a 30-day trial of the DriverLoader software to test it with your specific hardware, and also have the benefit of Linuxant's support personnel and extensive user community when diagnosing and resolving problems.**

Summary

This chapter explored one of the most liberating and empowering features of computing today, wireless networking. Wireless networking is not only convenient from the standpoint of having fewer cables strewn around your house but also enables people to do real work from their laptops, moving from conference room to conference room, office to office, and even city to city without having to plug their machines in and laboriously reconnect to and synchronize with a given network.

Bonus Chapter 1 explores how you can connect to other machines on your network. Whether you need a command-line login session or a remote graphical console to execute commands directly on a remote machine, or simply need to transfer files back and forth between multiple machines, most of the tools you'll need are installed by default on Ubuntu and Kubuntu desktop systems. Anything else that you may need is only a few clicks and a download away in the Ubuntu repositories.

Appendix A

What's on the DVD and CD-ROM?

This appendix provides you with information on the contents of the DVD and CD that accompany this book.

IN THIS APPENDIX

System Requirements

Using the DVD with Windows and Linux Systems

What's on the DVD and CD

Troubleshooting

System Requirements

Make sure that your computer meets the minimum system requirements listed in this section. If your computer doesn't match up to most of these requirements, you may have a problem using the contents of the DVD and CD.

For Linux or Windows Systems:

- PC with a Pentium processor running at 300 Mhz or faster; for best performance, we recommend a Pentium 4 or AMD equivalent running at 700 Mhz or better
- At least 384 MB of total RAM installed on your computer; for best performance, we recommend at least 512 MB
- Ethernet network interface card (NIC) or modem with a speed of at least 28,800 bps
- A DVD-ROM drive
- A VGA graphics card that is capable of at least 1024x768 resolution
- At least 4 GB of available disk space if installing Ubuntu Linux

If your system does not satisfy these requirements, you may still be able to install Ubuntu Linux on your system by either installing the Ubuntu Server edition or by downloading, burning, and using the Ubuntu Alternate Install CD. Both of these options are discussed in Chapter 3.

Using the DVD with Windows

To install Ubuntu Linux from the DVD to your hard drive without changing the partitions on your disk, follow these steps:

1. Insert the DVD into your computer's DVD-ROM drive.

 A window will appear, giving you the opportunity to install Ubuntu Linux on your Windows system by using the Wubi Installer.

 For more information using Wubi, see Chapter 2.

If you do not have autorun enabled or if the autorun window does not appear, follow the steps below to access the DVD.

1. Click Start Run.
2. In the dialog box that appears, type ***d:****umenu.exe*, where *d* is the letter of your DVD-ROM drive.
3. Choose the "Install inside Windows" option, as described in Chapter 2.

Using the DVD with Linux

NOTE **The DVD that is included with this book is a Ubuntu Linux distribution for 32-bit x86 systems, and will not work on 64-bit systems. To install Kubuntu on a 32-bit system, you must obtain and burn a copy of the Kubuntu CD for 32-bit systems. To install or test-drive Ubuntu or Kubuntu Linux on a 64-bit x86 system, you must obtain and burn a copy of the Ubuntu or Kubuntu CD for 64-bit systems. See the instructions on Chapter 2 for information on obtaining and burning a CD that contains the 32-bit/64-bit versions of Kubuntu or the 64-bit version of Ubuntu Linux.**

To install or test-drive Ubuntu Linux from the DVD to your hard drive, follow these steps:

1. Insert the DVD into your computer's DVD-ROM drive.
2. Reboot your computer system, booting from the DVD-ROM. If your system is not already configured to boot from a DVD before booting from the hard drive, you may need to modify your computer's BIOS settings in order to give precedence to booting from the DVD-ROM drive.
3. Press Return or wait 30 seconds to boot into Ubuntu Linux.
4. Experiment with Ubuntu Linux or follow the instructions in Chapter 2 to permanently install Ubuntu Linux on your computer system. If your system already contains an operating system that you want to preserve and has sufficient disk space for both that operating system and Ubuntu , the Ubuntu Linux installer will offer to repartition your disk without removing your existing data.
5. To remove the DVD from your DVD-ROM drive, select System Quit, and select Shut Down from the menu that displays. Remove the DVD-ROM from your computer's DVD-ROM drive when prompted to do so.

What's on the DVD

The DVD provides a version of Ubuntu Linux from Canonical, Ltd. that you can either test-drive or install on a 32-bit, x86-compatible computer system. You can install it on a new system or as an alternate operating system on an existing computer system. You can install it as an alternate operating system by either

reparititioning your disk to provide a partition where Ubuntu Linux can be installed directly, or as an alternate operating system on an existing Microsoft Windows system that has at least 4 GB of free space without repartitioning.

Shareware programs are fully functional, trial versions of copyrighted programs. If you like particular programs, register with their authors for a nominal fee and receive licenses, enhanced versions, and technical support.

Freeware programs are copyrighted games, applications, and utilities that are free for personal use. Unlike shareware, these programs do not require a fee or provide technical support.

GNU software is governed by its own license, which is included inside the folder of the GNU product. See the GNU license for more details.

Trial, demo, or evaluation versions are usually limited either by time or functionality (such as being unable to save projects). Some trial versions are very sensitive to system date changes. If you alter your computer's date, the programs will "time out" and will no longer be functional.

What's on the CD

The CD that accompanies this book provides twelve bonus chapters that simply didn't fit in the print edition of the book. These chapters are the following:

1. **Connecting to Other Systems:** Discusses graphical and command-line software such as SSH, VNC, and Remote Desktop that enables you to connect to other systems over a network.
2. **File Transfer and Sharing:** discusses the related topic of connecting to remote file servers to access files that are stored on them, and explains how to share files from your Ubuntu or Kubuntu system so that other systems on your network can quickly and easily access them.
3. **Software Development on Ubuntu and Kubuntu:** Explores how to install the software that you need to develop software on Ubuntu and Kubuntu systems, including compilers, related applications, graphical Integrated Development Environments, and so on. This chapter also includes a step-by-step examination of what it takes to build open source software from its source code, including how to identify additional packages that the software you're compiling requires.
4. **Managing Users, Groups, Authentication:** Explains how to manage users, groups, and system authentication (including Kerberos). This chapter describes how Ubuntu and Kubuntu systems use Pluggable Authentication Modules (PAMs) to provide a flexible and powerful authentication mechanism, how to customize the sudo command, and also discusses how to use Access Control Lists to provide more fine-grained control over who can access and modify certain files and directories.
5. **Managing System Policies and Security:** Compares and contrasts popular system policy enforcement software such as AppArmor and SELinux, and explains how to configure and use these packages.
6. **Backing Up and Restoring Files:** discusses various mechanisms for backing up and restoring files, starting with commands that enable you to create archives of specific files and directories, and then discussing mechanisms for backing up and synchronizing files and directories to other systems using the network.
7. **Using Virtual Machines and Emulators:** Virtualization and emulation are hot topics today, enaling you to run multiple operating systems on a single computer at the same time. This chapter provides an overview of emulation and the different types of virtualization that are available

today, and explains how to configure and use virtual machines using VirtualBox, VMWare Server, KVM, Xen, and QEMU. This is a really long chapter (80 pages or so), but there's a lot to cover to get you up and running with each of these technologies.

8. **Setting Up a Web Server:** Provides background information about the Web and Web servers, and explains how to install, configure, and troubleshoot the beloved Apache Web server.
9. **Setting Up a Mail Server:** Provides background information about mail servers, discusses the various mail servers that are available for Linux systems, and explains how to install and configure the Posfix mail server, including SPAM filtering and virus scanning.
10. **Setting Up a DHCP Server:** Explains DHCP, the Dynamic Host Control Protocol, and discusses how to install and configure DHCP servers manually or by using graphical tools.
11. **Setting Up a DNS Server:** Explains DNS, the Domain Name System, and discusses how to install, configure, and troubleshoot the popular BIND DNS server.
12. **Setting Up a Print Server:** Provides background information on printing from Linux and Unix systems, and explains how to install and configure CUPS, the Common Unix Printing System. This chapter also explains how to configure your Microsoft Windows and Mac OS X systems to print to you Linux print server.
13. **Setting Up an NFS Server:** Explores the history and different versions of NFS, the Network File System. Explains how to install and configure an NFS server on your Linux system (includign the latest NFS 4 servers), and also discusses how to mount NFS fileshares on client systems.
14. **Setting Up a Samba Server:** provides background information about Microsoft Windows file sharing, and explains how to install and configure Samba, a SMB/CIFS (Server message Block/ Common Internet File System) server that exports directories and printers from your Linux systems so that they can be used by other personal computer operating systems such as Microsoft Windows and Apple's Mac OS X.

Readers who also purchased the first edition of this book (thanks very much, BTW) will notice that bonus chapters 6 through 12 were actually Part IV of the first edition. These chapters have been updated and upgraded from the first edition, but since not everyone runs servers, I thought that they were good cantidates for moving to PDF-only chapters. My apologies if you disagree.

Troubleshooting

The DVD that is included with this book contains a complete version of Ubuntu Linux that you can either test-drive or install on your computer system. If you have problems installing or running Ubuntu Linux on a particular computer system, consult the Ubuntu Forums at `http://www.ubuntuforums.org/`.

Customer Care

If you have trouble with the CD-ROM, please call the Wiley Product Technical Support phone number at (800) 762-2974. Outside the United States, call 1(317) 572-3994. You can also contact Wiley Product Technical Support at `http://support.wiley.com`. John Wiley & Sons will provide technical support only for installation and other general quality control items. For technical support on the applications themselves, consult the program's vendor or author.

To place additional orders or to request information about other Wiley products, please call (877) 762-2974.

Index

SYMBOLS

(hash marks)
 /etc/apt/sources.list, 935
 bash shell, 566
comment (AppArmor), CD158
#include <file/containing/something> (AppArmor), CD158
#include statement (AppArmor), CD165
#ubuntu IRC channel, 21
$, 536
$TTL, CD327–329
* (asterisk), commands and, 537, 562
- (single dash), command-line options and, 540
- - (double dashes), command-line options and, 540
"." entries, -a option and, 538–539
".." entries, -a option and, 538–539
/ (forward slash)
 Linux file and directory names and, 104, 106
 ls command and, 537
/bin/ping application (AppArmor), CD160–161
/etc/apache2/sites-available/default file, CD282–283
/etc/apparmor.d directory, CD153
/etc/apt/sources.list file
 Adept Updater and, 980, 987
 repository components and, 916–917
 Software Sources tool and, 938
 system upgrades and, 951, 952
 Wicd and, 1095
/etc/bash.bashrc (bash configuration file), 565
/etc/bind/named.conf file, CD325, CD326
/etc/bin/named.conf.local file, CD326–327, CD332–333
/etc/dhcp3/dhcpd.conf file, CD311
/etc/exports file, CD363–365
/etc/fstab file, 111, 114
/etc/group file, 118, 119
/etc/hosts files, CD323, CD325
/etc/init.d/bind9 file, CD333
/etc/init.d/rcS command, 913
/etc/passwd file, 118
/etc/profile (bash configuration file), 565
/etc/selinux, CD154
/full/path/to/file/or/directory PERMISSIONS (AppArmor), CD159
/home Linux directory, 105
/opt Linux directory, 105
/proc/filesystems file, 115
/sbin/init, 911
/usr Linux directory, 105
/var Linux directory, 105
/var/log/messages log file, CD319
/var/www directory, CD282
? (question mark) wildcard, 562
@ sign, DNS servers and, CD328, CD329, CD330, CD331
[] (square brackets), commands and, 562, 563
\ (backslash) character, 181, 908
^ (caret symbol), 587
{ } (curly braces) wildcard, 562
~ symbol (bash), 565–566
~/.bash_profile (bash configuration file), 565

NUMBERS

3dchess package, 845
64-bit PC system installation, 74
64-bit Ubuntu Desktop CD, 29–30
802.11 standards, 1092–1093
802.11n standards, 1093

A

-a option, 538
AbiWord, 657

Accelerated Indirect GLX (AIGLX), 177
access
 to files and directories (AppArmor), CD159
 to printers, restricting, 1008–1009, 1010
Access Control Lists (ACLs), Linux
 activating filesystem support for, CD137
 basics, CD135–136
 installing commands, CD136–137
 types/format of, CD136
 working from command line, CD137–140
 working with graphical tools, CD140–142
access management. *See* AppArmor; SELinux
Accessibility desktop effects (KDE 4), 304–305
accounts (e-mail)
 default, 317
 defining for sending, 381–383
 FreePOPs!, 348–349
 IMAP, 379–380
 Modify Account dialog, 375–376
 POP, 324
 POP3, 376–379
accounts (IM)
 configuring in Kopete, 510–513
 configuring in Pidgin, 498–501
 connecting to in Kopete, 514–517
 connecting to in Pidgin, 504–505
 obtaining, 496–498
ACLs. *See* Access Control Lists (ACLs), Linux
ad hoc wireless networks, 1092
AdBlock extension (Firefox), 449
adding software. *See* software, adding/removing/updating
add-on software, building, CD78–79
add-ons (Firefox)
 adding, 444–450
 registering, 442
 viewing, 442–444
address blocks (Flash), 57
address books (e-mail)
 basics, 350–352
 synchronizing, 347
addresses
 DHCP, CD317, CD318
 e-mail, specifying, 317–318
 rewriting, CD297–298
adduser command, CD110
Adept Installer, 929–930
Adept Manager on KDE 3
 basics, 972–975
 software, managing, 976–979
Adept Manager on KDE 4
 basics, 980–982
 software, deleting, 987
 software, installing, 982–986
 software updates, searching for, 986–987
Adept Manager tools. *See* Adept Package Manager
Adept Notifier applet, 986
Adept Package Manager, 972–988. *See also* Adept Manager on KDE 3; Adept Manager on KDE 4; Adept Updater, on KDE 4
 vs. apt-get, 950
 basics, 972
 installing software and, CD64
 vs. Synaptic, 952
Adept Updater, on KDE 4, 980, 987–988
Adium X for Mac OS X, 496
Adobe Photoshop Express, 732
Advanced Linux Sound Architecture (ALSA), 750, 760, 778–779
AIGLX (Accelerated Indirect GLX), 177
AIM (AOL Instant Messenger), 497
AisleRiot solitaire, 826
Aitchison, Ron, CD322
Alacarte Menu Editor (GNOME), 134
alarms, for calendar entries, 362–363
Albitz, Paul, CD322
aliases, 568
alien package formats, 989, 991
Allman, Eric, CD290
ALSA (Advanced Linux Sound Architecture), 750, 760, 778–779

alternate install CDs, 19. *See also* install options from alternate CD
Amanda (The Advanced Maryland Automated Network Disk Archiver), CD181
Amarok
 adding Podcasts, 819
 basics, 814–816
 creating playlists with, 816–817
 digital audio players and (Kubuntu), 888–896
 playing audio files/Internet radio with, 814–819
 playing CDs with, 780–782
 playing/adding Internet radio stations, 818
Analog Clock (KDE 4), 293
Andreessen, Marc, 422
anonymous FTP servers, CD33
anti-aliasing, defined, 706
Anti-Spam Wizard, 411, 412
Anti-Virus Wizard, 411
Apache
 Apache 2 file locations, CD280–281
 basics, CD280
 configuring, CD282–284
 installing, CD280
 integrating backuppc with, CD189
 log files, CD285–286
Apache Server 2 Bible, 2nd Edition (Wiley, 2002), CD286
Apache2, CD101–105
APIs (Linux) sound, 773
AppArmor
 apparmor-profiles package, CD156–157
 basics, CD152–155
 emulating RBAC with, CD166–167
 installing and configuring, CD155–157
 profiles, CD157–161
 resources, CD167
 vs. SELinux, CD152–154
 utilities, CD161–162
AppArmor Administration Guide, CD160
Appearance
 desktop effects (KDE 4), 306–307
 panel (KDE 3), 235
Appearance Preferences application (Compiz), 178–180
Apple. *See* Mac OS X
applets, on KDE 3, 238–240
Application Launcher (KDE 4), 284, 285, 289, 293. *See also* K Menu
Application menu (GNOME), 130–131
application windows
 GNOME, 128–129
 KDE 3, 208–212
applications
 optimizing loading of, 919–920
 printing from, 1016, 1017
application-specific menus (KDE 4), 280
application-specific-variable definition (AppArmor), CD158
Applixware, 657
appointments (calendar entries), 361–364
AppSocket/HP JetDirect printer connection, 998, 1013
apt-cache utility
 for listing package information, 947
 searching for packages and, 946, 947
apt-file command, CD125
apt-file utility
 basics, 948
 for listing package contents, 949
apt-get, 950–954
 vs. aptitude, 954, 958–959
 basics, 931, 932
 build dependencies and, 953–954
 vs. dpkg, 954
 package source code, retrieving with, 952–953
 smart system upgrades, 951–952
 vs. Synaptic and Adept Manager tools, 950
 upgrading system with, 951
apt-get build-dep package command, CD67
apt-get command, CD66
apt-get source package command, CD67
aptitude, 954–959
 advantages of, 958–959

vs. apt-get, 954, 958–959
basics, 931, 932
vs. dselect utility, 954
installing BIND and, CD325
installing software with, 956–958
for listing package information, 947
search for emacs with, 945–946, 947
user interface, 955–956
arcade games, 831–835, 852–857
archive backups, CD174, CD175–176
archiving files, CD182–184
arguments, defined, 537
Ark graphical archiving tool, CD180
Arkeia's Network Backup, CD178
Armagetron arcade game, 857
ARPANET early history, CD321
"Art of Computer Programming", 606
aRts (analog real-time synthesizer sound server), 773
assistive technologies for GNOME, 155–156
asterisk(*), commands and, 537
Atlantik (game), 835, 847–850
Atomic Tanks arcade game, 852–853
attendees, adding to meetings, 364–365
audio (digital). *See* digital audio/video (Kubuntu); digital audio/video (Ubuntu)
audio file formats, converting, 820
audio files
playing with Amarok, 814–819
playing with Rhythmbox, 809–814
authentication
configuring Kerberos, CD142
defining e-mail security and, 383
options for retrieving e-mail, 380–381
Samba, configuring, CD376
authentication process (PAM), CD125–132
autoconf/automake applications, CD73
autodns-dhcp package, CD310
autofs package, CD366
automatically checking mail (Kontact), 396
automounting home directories (NFS), CD367
Avant Window Navigator (AWN), installing, 192–197
AVI files, working with, 490

B

Back In Time, CD198, CD199
background jobs, starting, 564–565
backgrounds
desktop, changing (GNOME), 152–153
desktop, changing (KDE 3), 247–248
in presentations, specifying, 697
backing up
archiving and restoring files, CD182–184
Backup Exec software, CD178
backup media, CD173
backup transports, defined, CD187
BackupEDGE software, CD178
backuppc, CD188
BackupPC package. *See* BackupPC package
bookmarks (Firefox), 439
defined, CD171
DVDs from command line, 792–793
e-mail files, 316–317
Linux time machines, CD198–201
to local, removable media, CD182
local directories, CD184–185
personal data for migration, 457–458
vs. RAID, CD172
reasons for, CD172–173
remote directories, CD185–187
software for Linux backups, CD177–182
types of backups, CD174–176
verifying and testing, CD176–177
what to backup, CD177
backports repository component, 934
backslashes (\), 908
BackupPC package, CD187–198
basics, CD181–182, CD187–188
identifying hosts, CD190–191
installing and configuring, CD188–190
restoring from backups in, CD195–198
rsyncd, CD191–193
SMB, CD193–194
starting backups in, CD194–195
Web interface, CD195, CD196, CD197
banners, adding, 1006, 1007, 1010
bare-metal restore, CD177

bash shell, 559–568
aliases, 568
command and filename completion, 561
command history, 560–561
configuration file, 565–566
environment variables, 566–567
input/output redirection, 563–564
job control, 564–565
pipes, 563–564
wildcards, 562–563
.bashrc configuration, 539
Basic Input/Output system (BIOS). *See* BIOS
Basic Service Set Identifier (BSSID), 1059
Battery Monitor (KDE 4), 294
Battle for Wesnoth arcade game, 857
bcm43xx-fwcutter, 1107
Beagle desktop search application, 162
Bell Labs, 543
Berkeley Mailbox (mbox). *See* mbox
Berners-Lee, Tim, CD278
Bernstein, Dan, CD289
Beryl, 177–178
Big MAC Attack, CD152
Binary Clock (KDE 4), 294
BIND (Berkeley Internet Name Daemon)
configuration, CD324
configuration files, CD325–327
errors, CD335–337
overview, CD322–325
resources, CD337
reverse lookup files and, CD332–333
BIOS
basics, 906
boot order, 30, 66
Boot Settings menu, 1030, 1031
booting from servers or alternate CDs and, 66–67
NetBIOS, CD372
bitmaps, defined, 706
BitTorrent, CD58–63
black and white scans, 1020, 1024
Blackjack, 827
blacklisting, defined, CD301
Blender, 742
blog (Ubuntu), 20
Blum, Richard, CD305
BMP, defined, 706
board games
Kubuntu, 835–837
Ubuntu, 827–829
Bogofilter, 411
Bonjour software, 499
The Book of Postfix: State-of-the-Art Message Transport (No Starch Press, 2005), CD305
bookmarks (Firefox), 434–441
backing up, 439
creating, 435–437
live, 428, 439–441
managing, 437–438
toolbar, 425
bookmarks, migrating to Linux, 481–489
basics, 481–482
exporting from Internet Explorer, 479, 482–484
exporting from Safari, 484–485
importing into Firefox, 486–487
importing into Konqueror, 487–489
boot loaders. *See* Grand Unified Boot Loader (GRUB)
boot process
BIOS Boot Device Priority settings, 1030
boot loader, 906–909
boot monitor or BIOS, 906
boot order, BIOS systems, 66
Bootchart and, 916–917
optimizing, 917–920
bootable Kubuntu USB stick, creating, 60
bootable Ubuntu USB stick, creating, 57–60
Bootchart, 916–917
booting
from alternate CDs, 66–67
from Desktop CDs, 30–34
servers for first time, 93
Ubuntu or Kubuntu, 45–47
BOOTP, DCHP and, CD308
Bourne, Stephen, 543
Brasero
burning CD/DVD images with, 798–799
burning CDs with, 793–798

bridged networking (virtual machines), CD246
bridge-utils package, CD227
browser settings, migrating to Linux
 basics, 475–476
 bookmarks. *See* bookmarks, migrating to Linux
 from Firefox to Firefox, 476–479
 from Internet Explorer/Safari to Firefox, 479–481
BRU Backup and Restore suite, CD178
BSSID (Basic Service Set Identifier), 1059
Buddy Lists (IM)
 adding in Kopete, 518–519
 adding in Pidgin, 505–507
 organizing in Kopete, 519–522
 organizing in Pidgin, 507
 viewing in Pidgin, 507–508
buffers, in emacs, 592–593
build dependencies, 953–954, 991–992
burning CDs
 with Brasero, 793–798
 with k3b, 801–805
burning data CDs/DVDs with Nautilus, 799–801
burning images
 CD/DVD, with Brasero, 798–799
 ISO, to CD/DVD with k3b, 805–806

C

C compiler, CD67–68
CA (certificate authority) certificates, 1060
cables, adding drives and, 1028
Cailliau, Robert, CD278
Calc
 basics, 672–674
 cell data type, specifying, 674–677
 files, installing, 670–671
 importing spreadsheets into, 677–680
 invoking functions in, 674
 starting, 671–672
 Ubuntu and Kubuntu and, 659–660
Calculator (KDE 4), 294
Calendar component (Evolution)
 basic configuration, 359–361
 meetings, 364–365
 online calendars, integrating, 366–367
 personal calendar entries, 361–364
 publishing calendars, 367–369
Calendar within Kontact
 basics, 416–419
 plug-ins, 393–394
camera interfaces (Kubuntu), 867
Canonical Ltd. paid support (Ubuntu), 23–24
Canonical's Launchpad, 15
Caplan, David, CD170
card games
 Kubuntu, 837–838
 Ubuntu, 826–827
case studies (Ubuntu examples), 48
case-sensitivity (vi), 574
cd (change directory) command, 552
CD/DVD Creator application, 799
CD/DVD drives, adding. *See* drives, adding
CD/DVD images, burning with Brasero, 798–799
cd'ing, 552
CD-Rs, for backups, CD173
CDs. *See also* install options from alternate CD; installing on special-purpose systems (Ubuntu/Kubuntu); Ubuntu Server Install CD options
 64-bit Ubuntu Desktop CD, 29–30
 booting from, 30–34
 burning ISO images to, 805–806
 burning with Brasero, 793–798
 burning with k3b, 801–805
 CD data, burning with Nautilus, 799–801
 copying with k3b, 807–808
 hard drive access from Desktop, 50–56
 installer CDs basics, 65
 installing Ubuntu/Kubuntu from, 34–47
 playing with Amarok, 780–782
 playing with Rhythmbox, 780
 ripping with k3b, 789–792
 ripping with Rhythmbox, 787–789
 setting blank preferences for (Kubuntu), 769–772
 setting blank preferences for (Ubuntu), 758–760
cell data, specifying
 Calc, 674–677
 Gnumeric, 684–686
 KSpread, 691–693

cells
data types or data formats, 673–674, 684, 691
defined, 660
formats, displaying, 691, 692
formatting, 667–668, 669
invoking functions in, 665
labeling, 666, 668
pasting values into, 664–665
CentOS, SELinux and, CD168
CERN (Conseil Européenne pour la Recherche Nucléaire), CD278
certificate authority (CA) certificates, 1060
cgoban game interface, 846
Chess, 828, 842–847
chgrp (change group) command, 558
Chicken of the VNC (client), CD12
chmod (change mode) command, 556–558
chown (change owner) command, 558
Christian Edition (Ubuntu), 14
chrootkit application, 1083
chsh command-line command, 544
CIFS (Common Internet File System), CD372
CinePaint, 742–743
ClamAV
basics, CD291, CD300
integrating with Postfix, CD303–305
clients
e-mail, defined, CD288
Evolution as, 313
IRC, 526
mounting NFS directories on, CD365–366
multi-protocol IM, 496
setting up Kerberos, CD147–148
clip art library, 706
Clock/Calendar applet (KDE 3), 208
clocks
Analog Clock (KDE 4), 293
system clock settings, 97
CMYK, defined, 706
CodeWeavers, 884
Cohen, Bram, CD58
color
aptitude screen and, 955–956
color depth, defined, 706
color palette (Krita), 725
color space, defined, 707
GIMP and, 714
pantone colors, 708
Color Lines (Windows game), 828
Color Picker (KDE 4), 294
Colton, Charles, CD198
columns
headings (Calc), 672
resizing, 667, 668
Comic Strip (KDE 4), 295
command line
accessing existing partitions from, 54–56
backing up DVDs from, 792–793
connecting to FTP servers from, CD40–43
connecting to Windows shares from, CD53
managing DHCP servers from, CD311–314
working ACLs from, CD137–140
command scripts, system run levels and, 912
command-line commands, 551–559
file and directory permissions, changing, 556–558
for Linux filesystems, 552
man, 559
vim, 580
for working with files and directories, 552–556
command-line interpreters, defined, 535
command-line mode (vi), 572
command-line systems, installing, 98
command-line tools, 535–568
aliases, 568
bash shell. *See* bash shell
command-line commands. *See* command-line commands
defined, 536
drives, adding and, 1038–1039
executing commands from, 536–540
information about commands, 541–542
Linux and, 535
to partition and format drives, 1039–1043
partitions, identifying and, 1038–1039
reasons for using, 536
shells basics, 542–544

starting shells and. *See* shells, starting
for wireless networking, 1102–1103
commands. *See also* command-line commands
completion of, 561
configuration commands (vim), 580
cursor movement commands (vi), 574–575
emacs, 586–587
executing from command line, 536–540
history of, 560–561
information about, 541–542
IRC, 526
keyboard commands (emacs), 591
man, basics, 559
recalling/re-executing, 560
running with KRunner, 288–289
screen movement commands (vi), 575
for searching and replacing (vi), 577–578
for testing DNS, CD334–335
commercial backup for Linux, CD178
commercial support for Ubuntu, 23–24
Common Internet File System (CIFS), CD372
Common UNIX Printing System. *See* CUPS (Common UNIX Printing System)
community support for Ubuntu/Kubuntu, 19–20
CompactFlash, 901
Compiz window manager
Appearance Preferences application, 178–180
Avant Window Navigator (AWN), installing, 192–197
CompizConfig Settings Manager, installing, 185–191
compositing basics, 176
configuring, 180–185
Emerald/Emerald Theme Manager, 197–199
history, 176–178
plug-ins, 192
special effects in, 178
using with KDE, 200–201
complain mode (AppArmor), CD155
components. *See* repositories and components
Compose Message window, 330
compositing
advanced graphics options, 302–303
basics, 176
enabling/disabling effects, 303–308
standard effects and shortcuts, 299–301
using in KWin, 297–298
Concurrent Versioning System (CVS), CD100
configuration variables (Postfix), CD295–297
configuration files
bash shell, 565–566
BIND, CD325–327
creating (DomainU), CD263–267
DHCP, CD311–314
Postfix, CD295
configuration models (Postfix), CD292–293
configuring. *See also* e-mail, configuring in Kontact
Apache, CD282–284
Apache2 for Subversion support, CD103–105
AppArmor, CD155–157
applets, 240
BackupPC package, CD188–190
Calendar component, 359–361
Calendar within Kontact, 390–395
common NFS options, CD361–362
Compiz with GNOME configuration tools, 180–185
configuration changes, undoing (KDE 3), 252
configuration files for individual PAMs, CD130–131
configuration files for network interfaces, 1077–1079
configure scripts, CD73
contacts (Kontact), 386–390
DNS configuration, manual, 1078
e-mail contacts, 349–353
encryption/authentication for network connections, 1098–1099
Firefox. *See* configuring Firefox
firewalls, 1083–1090
GStreamer framework and plug-ins, 775–776
hardware for networking, 1053–1054
IM accounts in Kopete, 510–513
IM accounts in Pidgin, 498–501
IMAP, 322
Kate, 600, 601
KDE 3 panels, 230–237

KDE 3 power management, 244–246
KDE monitor and display, 228
KEdit, 598
kernel server options (NFS), CD362–363
KMail, 373
Konqueror browser, 261–262
Kontact, 372–373
KWrite, 599–600
modems and dial-up connections, 1079
NFS servers, CD360
NTP on servers and clients, CD143–145
options for wireless networks, 1096, 1098
PAM configuration files, CD126–127
PAMS default configuration file, CD131–132
POP, for Gmail, 346
portable audio devices in Amarok, 888–889
Postfix, CD294–298
printers (Kubuntu), 1010–1012
printers (Ubuntu), 995–997
printers, existing (KDE 4), 1016–1017
printers manually (Ubuntu), 997–1002
PulseAudio, 776–779
Samba authentication, CD376
Samba server, CD374–379
screen resolution (GNOME), 140–142
screen resolution (KDE 4), 309–311
screensaver (GNOME), 147–148
screensaver (KDE 3), 242–243
sound on KDE 3, 760–762
sound on KDE 4, 763–765
sound on Kubuntu, 760–772
sound on Ubuntu, 750–760
spam/virus detection (Kontact), 411–414
Summary View (Kontact), 395–396
svnserve server, CD105–106
Synaptic preferences, 961–962
themes with System Settings application, 248
time zone (POP), 325–326
Tracker search application, 165–166
VPN connections (Ubuntu), 1062–1067
Wicd, 1100–1102
wired connections (Ubuntu), 1056–1057
wired/wireless connections on KDE 3, 1071–1074
wired/wireless connections on KDE 4, 1074–1076
wireless connections (Ubuntu), 1057–1061
Xen DomainU configuration files, CD263–267
configuring at installation (Ubuntu), 69–82
default login names, creating, 75, 76
encrypted volumes, 77, 78–79
geographic location, 69, 70
GRUB, installing, 80, 81–82
HTTP proxies, 77, 79
keyboard, specifying, 71
languages (spoken), 69–70
partitions, 73–75
server names, 71–72
server software for installation, 80, 81
specifying update mechanisms, 80
time zone, 72–73
user accounts, creating, 75, 76
user passwords, creating, 77–78
configuring Firefox
features support, 430–431
home page, 429–430
privacy settings, 431–434
connections
FreeNX connection package, CD27–28
krdc for terminal server, CD25–27
to other systems, CD3–4
rdesktop for terminal server, CD20–22
Remote Desktop Connection application (Kubuntu), CD16–19
SSH, enabling incoming, CD5–6
tsclient for terminal server, CD22–25
Conroy, Dave, 570
contacts. *See* e-mail contacts (Evolution); e-mail contacts (Kontact)
Content Scrambling System (css) package, 782
context-sensitive menus
Calc, 672, 673
defined (Firefox), 426
Draw, 734
GNOME, 132–134, 142–145
Gnumeric, 682
Impress, 699
Inkscape, 738

KDE 3, 215–217
KDE 4, 280–284
Krita, 725
mouse and (GNOME), 126
Nautilus file manager, 158
Writer, 620, 621
conversations, logging, 508–509, 523–524
convert utility, 718
converting audio file formats, 820
cookies (Firefox), 431–432, 433
Cool-Streams, 818
copying
CDs/DVDs with k3b, 807–808
directories, 552–556
in emacs, 592
files, 552–556
files over networks, 56
files/directories (Nautilus), 159–160
files/directories in Dolphin file manager, 257–258
in vi, 576–577
Coulouris, George, 570
Courier IMAP, CD299
Cox, Alan, 1084
cp command
backing up with, CD184–185
files and directories and, 553–555, 556
cpio utility, CD179
css (Content Scrambling System) package, 782, 785
Ctrl+P command, 560
Cube Gears plug-in (Compiz), 192
CUPS (Common UNIX Printing System)
information and references, CD352
log files, CD351
print servers, CD340–345
Web interface, accessing controls for, CD351–352
curly braces ({ }) wildcard, 562
currency in spreadsheets, 675–677, 683, 684–686, 693, 694
cursor movement commands (vi), 574–575
customizing
desktops (KDE 4), 308
digital device recognition (KDE 3), 868–871
digital device recognition (KDE 4), 871–873
digital device recognition (Ubuntu), 863–865
emacs, 585, 592–595
GIMP, 711
GNOME menus, 134–138
incoming message events (Kopete), 522–523
KDE 3, 223
key bindings, CD88–90
mouse behavior (GNOME), 138–140
mouse behavior (KDE 3), 224–227
panels (GNOME), 142–147
panels, contents of (KDE 3), 237–241
system sounds, 755–756, 766–768
vim, 580–581
cut and paste
in emacs, 592
in vi, 576–577
CVS (Concurrent Versioning System), CD100
Cyrus IMAP, CD299

D

DAC (Discretionary Access Control) systems, CD152
Dasher text entry system, 155
data. *See also* cell data, specifying
personal, backing up for migration, 457–458
recognized by KSpread, 691
Windows application data and migration, 489–491
data CDs/DVDs, burning with Nautilus, 799–801
data partitions, defined, 85
data types or data formats
Calc, 673–674
Gnumeric, 684
Kspread, 691, 692
spreadsheets, 674–675
D-Bus message bus system, 1053
DCIM directory, 901
dd backup utility, CD179
DEB (Debian) package format
converting to, 989
eliminating build dependencies with, 953
identifying dependencies and, 956
packages, CD72
Ubuntu and Kubuntu and, 926

debfoster package, 991
Debian Project, 10–11
Debian-based installers, 65–66, 94–99
deborphan utility, 991–992
debugging
 Firefox and, 450
 NFS operations, CD369–370
 power management (GNOME), 151–152
Decibel (KDE 4), 276
default e-mail accounts, 317
default permissions, 559
deleting
 directories in Dolphin file manager, 258
 files and directories, 552, 555–556
 files in Dolphin file manager, 258
 packages in Synaptic Package Manager, 968–970
 software. *See* software, adding/removing/updating
 text in vi, 575–576
Dent , Kyle, CD305
dependencies
 aptitude and, 958
 build dependencies, 953–954, 991–992
 Synaptic and, 961
Deskbar applet (GNOME), 145, 167–170
Desktop Access applet (KDE 3), 208
Desktop CD
 accessing existing partitions from, 50–56
 booting from, 30–34
 hard drive access from, 50–56
 installing from, 34–47
 Ubuntu, 18
Desktop Cube plug-in (Compiz), 192
Desktop Cube/Rotate Cube plug-in, 192
Desktop distribution (Ubuntu), 12
desktop publishing. *See* Scribus
desktop systems
 installing in OEM mode, 97–98
 installing in text mode, 96–97
desktops. *See also* GNOME desktop
 effects (KWin), 299–301, 303–308
 GNOME. *See* GNOME desktop
 icons for services/locations, CD30–32
 KDE 3. *See* KDE 3 (K Desktop Environment)
 KDE 4. *See* KDE 4 (K Desktop Environment)
 launcher, for xterm application, 549
 overview of (Linux), 123–125
development software, CD66
device drivers, defined, 993
Device Managers, 757–758, 1045–1046
Device Notifier (KDE 4), 290
df command, 555, 1038
DHCP (Dynamic Host Configuration Protocol), 1052–1053
DHCP servers, CD307–320
 dnsmasq, CD310
 managing from command line, CD311–314
 managing graphically, CD314–319
 servers, alternate, CD310
 servers, basics, CD308–309
 servers, installing, CD309–311
 troubleshooting, CD319–320
 udhcpd, CD310
 wide-dhcpv6-server, CD310
dhcpd3, CD310, CD311
Dia, 743
digiKam, 730, 731
Digikam application, 879–884
digital audio players
 Amarok and, 888–896
 overview, 884–885
 Rhythmbox and, 885–888
digital audio/video (Kubuntu)
 burning CDs, 801–805
 burning ISO images to CD/DVD, 805–806
 converting audio file formats, 820
 copying CDs/DVDs, 807–808
 installing/using soundKonverter, 822–824
 KInfoCenter application, 768–769
 multiple audio output devices, 763, 765
 playing audio CDs with Amarok, 780–782
 playing audio files and Internet radio on Kubuntu, 814–819
 playing audio files and Internet radio using Rhythmbox, 809–814
 playing DVDs with Dragon, 786–787
 playing DVDs with Kaffeine, 784–786

proprietary formats, 774
setting blank CD/DVD preferences, 769–772
system sound devices, specifying, 760–764
system sound levels, setting, 765–766
system sounds, testing and customizing, 766–768
digital audio/video (Ubuntu)
backing up DVDs, 792–793
burning CD/DVD images with Brasero, 798–799
burning CDs with Brasero, 793–798
burning data CDs/DVDs with Nautilus, 799–801
Device Managers, 757–758
installing software for encrypted DVDs, 782
installing/using Sound Converter, 820–822
playing audio CDs, 780
playing DVDs, 783–784
PulseAudio, configuring, 776–779
ripping audio CDs with k3b, 789–792
ripping audio CDs with Rhythmbox, 787–789
setting blank CD/DVD preferences, 758–760
setting system sound levels, 753–755
sound devices, specifying and testing, 751–753
system sounds, testing and customizing, 755–756
terminology, 745–750
digital cameras
on Kubuntu, 879–884
on Ubuntu, 873–878
Digital Clock, 290, 295
digital devices
recognition on Kubuntu, 865–873
recognition on Ubuntu, 860–865
digital graphics. *See* graphics
Diller, Antoni, 615
directories
backing up, CD184–187
basics (Linux), 103–106
changing permissions, 556–558
common, 85
copying/moving/renaming/deleting, 257–258, 552–556
creating, 258
with critical personal data (Windows), 458
defined, 104
Evolution, 317
for mail storage, 461
managing in Nautilus, 159–162
NFS, mounting on clients, CD365–366
relationship to files, 104
selecting multiple, 257
sharing, CD379–380
Discretionary Access Control (DAC) systems, CD152
Disk Manager, 1032
disks
adding. *See* logical volumes (Linux)
adding internal, 1027
defining virtual, CD248
partitioning, manually, 42–45
partitioning and, 106–107
preparing, 106–107
Xen configuration files for, CD265
DISPLAY environment variable, 124
distributed filesystems, CD356
distributions (Linux), 5–6
DistroWatch, 24
dithering, defined, 707
DKMS (Dynamic Kernel Module Support), 994
dlocate utility, 949
dmesg command, 1038
DNS
manual configuration, 1078
origins of, CD307
DNS and BIND (O'Reilly), CD322
DNS servers, CD321–338
BIND, CD325–327
installing, CD325
overview, CD322–325
resources, CD337
restarting and testing, CD333–335
reverse lookup files, CD331–333
troubleshooting, CD335–337
zone files, CD327–331
zone files, incorporating with BIND, CD332–333
dnsmasq, CD310
DOC files, 491
docking and undocking dialogs, 712–713
Document Converter Wizard
Calc and, 670

importing from other word processors and, 631
importing old-format SXW files and, 617
importing presentations with, 702
importing spreadsheets with, 678–680
Document Wizard (Karbon14), 739–740
documentation (Ubuntu), 22–23
documents, creating and publishing, 605–657
AbiWord, 657
Applixware, 657
basics, 605–606
markup languages. *See* LaTeX; TeX
publishing with Scribus. *See* Scribus
word processing. *See* KWord; OpenOffice.org
Dolphin
accessing existing partitions with, 52–54
basic operations, 257–258
connecting to FTP servers in, CD38–40
connecting to Windows shares in, CD52–53
overview, 253–256
"The Domain Names Plan and Schedule", CD321
"Domain Names — Concepts and Facilities", CD321
"Domain Names — Implementation and Specification", CD321
domains
accepting mail for, CD298
identifying (Postfix), CD297
DomainU
building paravirtualized, CD267–272
configuration files, creating/configuring, CD263–267
hardware virtual machine, creating, CD271–272
Doom arcade game, 853–854
downloading. *See also* Web sites for downloading
VMware server, CD235–242
dpkg
vs. apt-get, 954
basics, 931, 932
command's --search option, 948
listing package contents with, 947–948
listing packages with, 943–946
Dragon player, playing DVDs with, 786–787
Drawer applet (GNOME), 145
drive bays, defined, 1027
DriverLoader, 1108
drivers
adding hardware and, 993–994
device drivers, defined, 993
selecting for printers, 1006, 1007
drives, 1033
drives, adding
command-line tools to identify, 1038–1039
command-line utilities to partition and format, 1039–1043
EIDE/ATA drives, 1027–1028
formatting with GNOME Partition Editor, 1033–1037
hard drive names, 1029
locating, partitioning and formatting new drives, 1032–1033
SATA drives, 1029
SCSI drives, 1030
troubleshooting boot problems, 1030–1032
Dropbear, CD162
dselect utility
vs. aptitude, 954
basics, 931, 932
dual-band cards, 1093
dual-boot systems
booting and, 93
booting Ubuntu/Kubuntu on, 46
GRUB boot loader and, 80, 81
migration and, 489–491
partitions and, 74
dump/restore utility, CD179–180
DVD-Rs, for backups, CD173
DVDs
backing up from command line, 792–793
burning ISO images to, 805–806
copying, 807–808
data, burning, 799–801
installing software for encrypted, 782
playing with Dragon on KDE 4, 786–787
playing with Kaffeine on KDE 3, 784–786
playing with Totem, 783–784
setting blank preferences for (Kubuntu), 769–772
setting blank preferences for (Ubuntu), 758–760

DVI files, formatting and previewing, 610–611
dviutils package, 611
Dynamic Host Configuration Protocol (DHCP), 1052–1053. *See also* DHCP servers
Dynamic IP addresses, 1053
Dynamic Kernel Module Support (DKMS), 994
Dynamic WEP (802.1x), 1060

E

EAP (Extensible Authentication Protocol), 1060
eboard chess display application, 844
echo command, 566–567
Eclipse IDE
 for C application development, CD83
 creating new source projects in, CD85–93
 defined, CD81
 importing existing source project into, CD93–99
 information references, CD99
 installing, CD82–83
 starting, CD83–84
 terminology, CD85
ed, 570, 602
editing
 documents in KWord, 637–638
 e-mail contact information, 354–356
 ID3 tag, 885
 iPod playlists, 898–900
 photos, 730–733, 876–877
 playlists, 894–895, 898–900
 text files. *See* text files, editing
Editing MACroS. *See* emacs
Edubuntu, 7, 13
eiciel, CD136–137, CD140
EIDE/ATA (Enhanced Integrated Drive Electronics/Advanced Technology Attachment) drives, adding, 1027–1028
emacs, 583–595
 basics, 583
 commands, 586–587
 customizing, 585, 592–595
 cutting, copying and pasting in, 592
 modeless editors and, 571
 moving around in, 591–592
 multiple windows and buffers in, 592–593
 resources, 595
 searching for, 945–946
 starting and exiting, 588–591
 startup, 583–584
 terminology, 587–588
 text, searching and replacing in, 592
 vs. vi, 571
 XEmacs, 585–586
e-mail, configuring in Kontact
 accounts, defining for sending, 381–383
 basics, 373–377
 identity, defining, 384–386
 IMAP accounts for receiving, 379–380
 POP3 accounts for receiving, 377–379
 security, defining, 380–381, 383–384
e-mail, free and Web-based
 Gmail, 342–347
 providers, 347–349
e-mail, migrating to Linux, 458–475
 basics, 458–459
 manually, 473–475
 via Thunderbird. *See* Thunderbird, migrating e-mail via
e-mail, sending and receiving, 326–338
 getting started, 326–329
 mail folders, 334–335
 preferences for sending and viewing, 331–333
 receiving, set-up for, 315–319
 receiving vs. retrieving, CD288, 316
 search folders, 336–338
 sending, 323–325, 330–331
 sending (Kontact), 404–408
 sending, set-up for, 323–325
 undeleting messages, 334
e-mail basics, CD287, CD288. *See also* mail servers
e-mail contacts (Evolution)
 configuring support for, 349–353
 contact information, adding/editing, 354–356
 mail lists, defining, 357–359
 vCards, 356
e-mail contacts (Kontact)
 configuring, 386–390

creating, 414–416
managing, 414–416
embedded devices, 859–860
Emerald window decorator, 197–199
emulated machines, managing, CD209
emulating RBAC, CD166–167
emulation software for Linux, CD208–209
encryption
defining e-mail security and, 383
encrypted DVDs, software for, 782
encrypted volumes at installation, 77, 78–79
enforce mode (AppArmor), CD155
environment variables, 566–567
errors
BIND, CD335–337
DNS servers, CD335–336
error message about repositories, 80
zone files, CD335–337
Escape sequences, defined, 587
ESD (Enlightened Sound Daemon), 773, CD75
esvn client, CD101
Ethernet networking, 1050
events, scheduling on Calendar, 360–361
Evolution, 313–370
Calendar component. *See* Calendar component (Evolution)
contacts, managing. *See* e-mail contacts (Evolution)
Evolution Setup Assistant. *See* Evolution Setup Assistant
filtering incoming mail, 338–342
free and Web-based mail. *See* e-mail, free and Web-based
fundamentals, 313–314
importing mail into, 466–469
launching, 314–315
migrating from older versions, 317
resources, 369–370
sending and receiving. *See* e-mail, sending and receiving
switching to, 463
vs. Thunderbird, 461
toolbar, 328–329
Evolution Setup Assistant
IMAP, 319–320, 322
incoming mail server parameters, 320–322
migrating from older versions, 317
POP, 319–322
Restore Evolution, 316–317
sending e-mail, 323–325
setting up for receiving e-mail, 315–319
time zone, 325–326
ex, 570, 602
Exchange Server, 314
Exim, CD288–289
exiting
emacs, 588–591
Kopete IM client, 524–525
Pidgin IM client, 509
vi, 572
exported directories
defining in /etc/exports (NFS), CD363–365
pseudo-filesystem (NFS) and, CD363
exported filesystems, verifying (NFS), CD368–369
exporting
bookmarks from Internet Explorer, 479, 482–484
bookmarks from Safari, 484–485
documents, 645
to online photo services, 877–878
ext2 filesystem, 109
ext3 filesystem, 110
Extended partitions, 91
Extensible Authentication Protocol (EAP), 1060
extension points, defined, 1092
extensions. *See* add-ons (Firefox)

F

-F option, 537
favorites, exporting from Internet Explorer, 482–484. *See also* bookmarks (Firefox)
fdisk command, 112, 1039–1042
fdisk utility, -l (list) option, 111
Fedora, SELinux and, CD168
Feldman, Stu, CD68
FFMPEG project, 803
fields

in /etc/fstab entries, 114–116
first fields, package entries, 943–944
SOA records, CD327–329
Fifteen Puzzle (KDE 4), 295
file extensions
apt-get source utility, 953
emacs modes, 594
OpenOffice.org Writer, 617
when saving GIMP files, 716
file handles, defined, CD358
file managers
Dolphin, 257–258
using Konqueror browser as, 262–264
file menus, 104
file sharing
ACLs for. *See* Access Control Lists (ACLs), Linux
FTP usage. *See* FTP (File Transfer Protocol) servers
overview, CD30
peer-to-peer (Ubuntu). *See* P2P file sharing (Ubuntu)
Windows shares. *See* Windows shares
File Transfer Protocol servers. *See* FTP (File Transfer Protocol) servers
File Watcher (KDE 4), 295
filename completion, 561
files. *See also* text files, editing
archiving and restoring, CD182–184
backing up to removable media, CD182
changing permissions, 556–558
checking out from Subversion repository, CD106
copying over networks, 56
copying/moving in Dolphin, 257–258
copying/moving/renaming/deleting, 552–556
defined, 103
deleting in Dolphin, 258
determining packages for missing, 949–950
file locations (Apache 2), CD280–281
installing, 670–671, 695–696
Linux, 103–106
managing in Nautilus, 159–162
relationship to directories, 104
renaming in Dolphin, 258
selecting multiple in Dolphin, 257
working with various types of, 490–491
filesystems
commands for, 552
defined, 107
initial RAM, 909–911
Nautilus, 159
support for ACLs, CD137
verifying exported (NFS), CD356
filesystems (Linux), 106–111
commands for, 552
disks, partitions and mount points, 106–107
local, 109–110
logical volumes, 107–109
mounting, 111–114
mounting at boot time, 114–117
network filesystems, 110–111
removable media filesystems, mounting, 117
FileZilla FTP client, CD43–44
filter commands, defined, 572
filtering, packet, 1083
filters
e-mail, 338–341, 409–411
KDE 3 Adept Manager, 974–975
Firefox, 421–453
add-ons, adding, 444–450
add-ons, registering, 442
add-ons, viewing, 442–444
bookmarks. *See* bookmarks (Firefox)
compared to other browsers, 421, 422, 427, 475
configuring. *See* configuring Firefox
Firefox Import Wizard, 486
history of, 421–422
importing bookmarks into, 486–487
migrating to, 476–479
migrating to, from Internet Explorer/Safari, 479–481
plug-ins, adding, 450–453
starting, 422–424
svn URLs in, CD106
terms, Internet Explorer and, 476
themes, adding, 444–450
themes, viewing, 442–444
user interface, 424–428
firewalls, installing and configuring, 1083–1090
first fields, package entries, 943–944

first-stage boot loaders, 906
Five or More, 828
FLAC (Free Lossless Audio Codec) package, CD79–80
Flash drives
 overview, 57
 reformatting USB, 59
Fluendo, 776
Fluxbuntu, 13
FlyBack, CD198, CD199
folders
 e-mail, 334–335, 375–376, 408–409
 search folders (e-mail), 336–338
formats
 Calc spreadsheet imports, 677–678
 converting audio files, 820
 graphics, 736
 graphics, scanning, 1025–1026
 proprietary audio/video, 774
 supported by GIMP, 719–721
 supported by Krita, 722–723
 which Gnumeric can open, 688
 which KSpread Calc open, 694–695
formatting
 cells, 667–668, 669
 drives, basics, 1032–1033
 drives, using command-line utilities, 1039–1043
 drives, using GNOME Partition Editor, 1033–1037
 DVI files, 610–611
forums (Ubuntu), 20–21
forwarding e-mail, 398
Foundation (Ubuntu), 10
Four-in-a-Row, 829
frames
 defined, 634
 editing options for (Scribus), 648, 649, 650
 KWord and, 638–644
 text frames, 653–656
framesets, defined, 634
free e-mail. *See* Gmail, configuring; Gmail, setup for reading
Free Internet Chess Server, 843
Free Lossless Audio Codec (FLAC) package, CD79–80
Free Software Foundation (FSF), 3, CD67
free/busy information
 configuring, 392–393
 publishing, 367–369
Freecell solitaire, 827
Free-java-sdk IDE, CD82
freenode IRC network, 530
FreeNX connection package, CD27–28
FreePOPs!
 basics, 347
 installing and using, 348
 POP accounts for, 348–349
FreeRADIUS server, 1060
The Fridge, 22
Frozen Bubble arcade game, 853–855
fsck utility
 fscking filesystems, 109
 time needed for, 109
FSF (Free Software Foundation), 3, CD67
F-Spot application, 873–878
F-Spot Photo Manager, 730, 731
FTP (File Transfer Protocol) servers
 basics, CD32
 connecting from command line, CD40–43
 connecting from GNOME desktop, CD32–35
 connecting from KDE desktop, CD35–40
 connecting in Dolphin, CD38–40
 connecting in Konqueror, CD36–38
 connecting with graphical clients, CD43–44
 setting up, CD45
functions
 defining (emacs), 594–595
 invoking in Calc, 674
 invoking in Gnumeric, 682–683
 invoking in KSpread, 690–691
Fuzzy Clock, 295

G

Gadu-Gadu, 499
games
 arcade games (Kubuntu), 831–835
 arcade games (Linux), 852–857
 board games (Kubuntu), 835–837

board games (Ubuntu), 827–829
card games (Kubuntu), 837–838
card games (Ubuntu), 826–827
Chess games, 842–847
for kids (Kubuntu), 838–839
on Kubuntu, 830–842
menu (Kubuntu), 841–842
Monopoly clones, 847–851
Scrabble clones, 851–852
tactics and strategy games (Kubuntu), 839–841
on Ubuntu, 825–830
video games (Ubuntu), 829–830
GCC (GNU Compiler Collection) compilers, CD66–68
gconf-editor (graphical application), 181–182
gconftool-2 command-line utility, 181–182
gdhcpd package, CD314–315
GDM (Graphical Display Manager), 918
gedit, 595–597
geographic location, configuring at installation, 69, 70
Get Ubuntu Extensions link, 447
Geubuntu, 13
gFTP client, CD43
GIF, defined, 707
GIMP (GNU Image Manipulation program), 709–721
basics, 709
color and, 714
converting and reducing images, 718–719
formats supported by, 719–721
quick tour, 711–713
resources, 721
starting, 710–711
taking screenshots, 714–718
GIMP Toolkit (GTK), 680
Gmail, configuring, 374
Gmail, setup for reading
basics, 343
transport protocol, selecting, 343–345
using IMAP, 346
using POP, 345
GNOME
configuring Compiz with, 180–185
connecting to FTP servers from, CD32–35
connecting to Windows shares from, CD48–50
GNOME DHCP configuration utility, CD315
gnome-screenshot command, 714
vs. KDE, 203–205
menus. *See* menus (GNOME)
mouse actions in, 205
GNOME desktop
application windows, 128–129
assistive technologies, 155–156
backgrounds, changing, 152–153
Deskbar applet, 167–170
keyboard shortcuts, 156–157
mouse behavior, customizing, 138–140
mouse usage, 126
Nautilus file manager, 157–162
overview, 126–127
panels, customizing, 142–147
power management, configuring, 148–152
screen resolution, configuring, 140–142
screensaver, configuring, 147–148
themes, 153–155
Tracker desktop search application, 162–166
window managers, 170–174
GNOME Device Manager, 757
GNOME Format, 1032
GNOME Partition Editor, formatting drives with, 1033–1037
GNOME terminal application, 545–546
GNOME vim (gvim), 582
gnome-app-install, 931
gnome-nettool package, 1080
gnome-ppp package, 1079
Gnometris, 830
GNU Chess, 828, 843–844
GNU Compiler Collection compilers (GCC), CD66–68
GNU Image Manipulation program. *See* GIMP (GNU Image Manipulation program)
GNU Project, 3
gnubg (GNU backgammon), 835
gnugo Go-playing engine, 846
GNU/Linux distribution, 6
Gnumeric spreadsheets, 680–687
basics, 680, 681–682
cell data type, specifying, 684–686

Help, 682
importing spreadsheets into, 686–687
invoking functions in, 682–683
starting, 681
GNUPaint, 743
Go (game), 845–847
Gobuntu, 7, 13–14
Google
Picasa photo management software, 730, 732
Talk IM application, 531
Gosling, James, 570
Grand Unified Boot Loader (GRUB), 80, 81–82, 906–907, 1031–1032
granularity, setting (Calendar), 360–361
graphical clients, connecting to FTP servers with, CD43–44
graphical Disks utility (Ubuntu), 1032
Graphical Display Manager (GDM), 918
graphical environments for Linux, 123–125
graphical installers vs. Debian-based, 65–66
graphical interfaces, supporting tools and, 67
graphical tools
Ark graphical archiving tool, CD180
Software Sources tool, 937–938
TeX and LaTeX and, 611–614
vector graphics tools for Linux, 735–736
working ACLs with, CD140–142
graphical versions of vi, 582
graphically managing DHCP servers, CD314–319
graphics, 705–743
applications, 742–743
Blender, 742
boot process and, 917
CinePaint, 742–743
for frames, 650–653
GIMP. *See* GIMP (GNU Image Manipulation program)
Inkscape, 736–738
Karbon14, 739–742
Krita. *See* Krita
OpenOffice.org Draw, 733–735
options, configuring advanced, 302–303
photo editing, 730–733
scanning formats, 1025–1026
terminology, 706–709
vector graphics tools for Linux, 735–736
GraphicsMagick package, 723
grayscale, defined, 707
greylisting, CD300–301
groups, 118–119
groups, system (KDE 3)
creating new, CD119–120
management overview, CD116–117
managing existing, CD120
groups, system (KDE 4)
creating new, CD123–124
managing existing, CD124
groups, system (Ubuntu)
creating new, CD114–115
management overview, CD110–111
managing existing, CD115–116
Groupwise software, 499
GRUB (Grand Unified Boot Loader), 80, 81–82, 906–907, 1031–1032
GStreamer framework/plug-ins, configuring, 775–776
gstreamer0.10-plugins-ugly-multiverse package, 789
gthumb, 873
gtk-gftp GNOME FTP client, CD43
gtk-gnutella, installing/using, CD56–58
gtkpod application for iPods, 896–900
gtk-window-decorator, 197
Guest Additions (VirtualBox), installing, CD224
Guest OS virtualization, CD205, CD208
gufw graphical configuration utility, 1085–1090

H

HAL (Hardware Availability Layer), 1053
Haley, Chuck, 570
halftones, defined, 707
hard disks. *See* hard drives
hard drives
access from Desktop CD, 50–56
changes to BIOS hard disk, 1030
creating virtual for KVM, CD254–255

names, 1029
portable, 1027
Hardening Apache (Apress, 2004), CD286
Hards, Brad, CD9
hardware. *See also specific* hardware
adding, 993–994
configuring for networking, 1053–1054
examining and troubleshooting, 1045–1046
requirements (Ubuntu/Kubuntu), 17
support for virtual machines, CD207
virtualization, CD206
Hardware Availability Layer (HAL), 1053
hardware virtual machine DomainU, creating, CD271–272
hash marks (#)
/etc/apt/sources.list, 935
bash shell, 566
Hazel. Philip, CD288
Heimdal client/server implementation, CD142
Help
Calc, 672, 673–674
Draw, 733, 734–735
Firefox, 429
Gnumeric, 682
Impress, 699
Inkscape, 737
Karbon14, 741–742
Krita, 725, 726
KSpread, 690
KWord, 637
ndiswrapper, 1108
Synaptic, 959–961
Writer, 620, 621
-help option, commands and, 540–541
The Hessling Editor (THE), 603
hibernate mode, 150
hiding panels automatically (KDE 3), 232
Hildebrandt, Ralf, CD305
home directories
automounting (NFS), CD367
sharing (NFS), CD356
sharing using Samba, CD376–377
home pages
Firefox, 423–424, 429–430
Gnumeric, 682, 683
importing from Internet Explorer, 480–481
hostnames, CD323, CD324
host-only networking (virtual machines), CD245
hosts
defined, CD278
identifying, CD190–191, CD297
mail, www, and ftp hosts, CD331
HOSTS.TXT file, CD321
Hotmail, configuring, 374
HSV, defined, 707
HTML. *See* HyperText Markup Language (HTML)
hub, defined, 1051
HyperText Markup Language (HTML)
Kontact preferences, setting and, 401–404
Message Preview pane, 396
sending and viewing e-mail and, 331–333
HyperText Transfer Protocol (HTTP)
basics, CD278
proxies, 77, 79

I

I Seek You (ICQ) software, 495
i18n (internationalization), defined, 8
Iagno, 828
Ibanez, Roger Ferrer, CD140
icons, desktop, CD 30–32
ICQ ("I Seek You") software, 495, 497
ID3 tag editing, 885
IDE drives, problems after adding to SATA systems, 1031
identity (Kontact), defining, 384–386
IDEs (Integrated Development Environments)
basics, CD80
Eclipse. *See* Eclipse IDE
for Linux, CD81–82
Image Wizard (Krita), 723–724, 727
ImageMagick, 718, 723
images
background (KDE 3), 247–248

converting and reducing (GIMP), 718–719
loading (Firefox), 430–431
IMAP
accounts for receiving e-mail, 379–380
configuring options, 322
Gmail, accessing with, 346
MDAs and MUAs and, CD288
migrating e-mail and, 459
vs. POP, CD294, 319–320
vs. POP3, 374
Import Wizard (Evolution), 466, 467–469
Import/Export Wizard (Internet Explorer), 482–484
importing
documents, 630–633, 645
presentations (Impress), 702–703
spreadsheets into Calc, 677–680
spreadsheets into Gnumeric, 686–687
spreadsheets into KSpread, 693–695
spreadsheets with Document Converter Wizard, 678–680
Impress, 695–702
basics, 698–699
Help, 699–700
Impress Presentation Wizard, 696–698
installing files for, 695–696
presentations, creating, 700–702
presentations, importing, 702–703
starting, 696–698
Ubuntu and Kubuntu and, 659–660
incremental backups, CD174, CD175–176
inetd daemon, CD105
info command, 542
init (initialization) process, 911–916
basics, 911
run levels, 912
SysVInit system initialization, 913
Upstart system initialization, 913–916
init packages, 911
initramfs (initial RAM filesystem), 909–911
initrd (initial RAM disk), 909–911
Inkscape, 736–738
input redirection, 563–564
install CDs. *See* install options from alternate CD; installing on special-purpose systems (Ubuntu/Kubuntu); Ubuntu Server Install CD options
install options from alternate CD
basics, 94–96
installing command-line system, 98
LTSP server, 99
OEM mode, 97–98
text-mode installer, 96–97
installing. *See also* configuring at installation (Ubuntu)
Apache, CD280
AppArmor, CD155–157
Avant Window Navigator (AWN), 192–197
BackupPC package, CD188–190
CompizConfig Settings Manager, 185–191
DHCP servers, CD309–311
DNS servers, CD325
Eclipse, CD82–83
Emerald, 197
files for OpenOffice.org Writer, 617–618
firewalls, 1083–1090
FreePOPs!, 347–349
GCC compilers, CD66
Guest Additions (VirtualBox), CD224
KDE on Ubuntu, 271
KDE-related packages, 200
Kerberos client and server commands, CD142–143
Kubuntu. *See* installing on special-purpose systems (Ubuntu/Kubuntu)
KVM, CD253–254
KWord, 634
LaTeX, 606–607
NFS servers, CD360
operating systems in virtual machines, CD220–223, CD255–256
packages in Synpatic, 964–968
Postfix, CD291–294
problems with boot/installation (virtual machines), CD233–234
requirements for (Ubuntu), 16–18
Samba server, CD373–374

Scribus, 646–647
SELinux, CD168–169
server software for, 80
software for encrypted DVDs, 782
software with Adept Manager, 977, 982–986
Sound Converter, 820–822
soundKonverter, 822–824
Subversion, CD100–101
TeX, 606–607
themes with Kthememanager application, 250–251
Thunderbird, 462
types of installers, 65–66
Ubuntu. *See* installing on special-purpose systems (Ubuntu/Kubuntu); installing Ubuntu/Kubuntu
VirtualBox, CD212–213
VMware server, CD235–242
Wicd, 1095
Xen technology, CD263
installing on special-purpose systems (Ubuntu/Kubuntu), 65–99
alternate CD install options. *See* install options from alternate CD
booting from alternate CDs, 66–67
install CDs basics, 65–66
install CDs, obtaining, 66
Ubuntu Server Install CD options. *See* Ubuntu Server Install CD options
installing Ubuntu/Kubuntu
64-bit Ubuntu Desktop CD, 29–30
bootable Kubuntu USB stick, creating, 60
bootable Ubuntu USB stick, creating, 57–60
booting, 45–47
booting from Desktop CDs, 30–34
Desktop CD persistence, 56–60
from Desktop CDs, 34–47
examples folder, 47–50
hard drive access from Desktop CD, 50–56
installation requirements, 16–18
migration, and installing, 489–491
partitioning disks manually, 42–45
selecting version of, 28–29
on Windows with Wubi, 60–63
instant messages (IM)
accounts, obtaining, 496–498
communication problems, resolving, 532–533
IM application packages, 531–532
IRC. *See* IRC (Internet Relay Chat)
overview, 493–495
on Ubuntu with Kopete. *See* Kopete IM client
on Ubuntu with Pidgin. *See* Pidgin IM client
Integrated Development Environments. *See* IDEs (Integrated Development Environments)
International Go federation, 847
internationalization
Calc and, 670–671
defined (i18n), 8
Impress and, 695
Internet, history of, CD321–322
Internet Explorer
exporting bookmarks from, 482–484
Firefox terms and, 476
migrating from, to Firefox, 479–481
Internet Printing Protocol (ipp) printer connection, 998, 1013
Internet radio
playing with Amarok, 814–819
playing with Rhythmbox, 809–814
Internet radio stations
playing/adding with Amarok, 818
playing/adding with Rhythmbox, 811–813
Internet Relay Chat. *See* IRC (Internet Relay Chat)
Internet Scrabble Club, 851
Internet service managers, CD185
Internet service providers (ISPs), e-mail and, 316
invitations to Calendar events, 419
IP addresses, 1052
iPods. *See also* digital audio players
adding music to, 897–898
gtkpod for working with, 896–900
playlists, editing, 898–900
ipp (Internet Printing Protocol) printer connection, 998

iptables interface, 1084–1085
IRC (Internet Relay Chat)
 #ubuntu IRC channel, 21
 clients, 526
 overview, 525–526
 sources of information about, 531
 using in Kopete, 528–530
 using in Pidgin, 526–528
ISO images, burning to CD/DVD, 805–806
iTunes, running on Linux, 884
iwconfig command, 1102
iwlist command, 1102–1103

J

Jabber IM platform, 497
JACK Audio Connection Kit, 773
jaggies, defined, 707
Japanese Shogi (chess game), 845
Java, enabling (Firefox), 430–431
JavaScript, enabling (Firefox), 430–431
Jduplicate (Scrabble clone), 851
jedit, 603
JEOS (Just Enough Operating System)
 installation configuration, 68–69
 Ubuntu Server Edition, 16
JFS filesystem, 110
job control, 564–565
job files, system initialization and, 913–914
jobs, Upstart and, 914
journaling local filesystems, 109–110
Joy, Bill, 570, 579
JPEGs
 defined, 707
 reducing quality when saving, 718–719
JuK media player, 784
jumpers, 1027, 1028
junk e-mail
 checking automatically, 341–342
 filtering, 338–341
 folder for, 335

K

-k (keyword search option), 559
K Menu
 applet, 207–208, 238
 Applications System Printing application, 1012, 1016
 basics (KDE 3), 213–214
 basics (KDE 4), 284–288
k3b
 burning CDs with, 801–805
 copying CDs/DVDs with, 807–808
 ripping CDs with, 789–792
Kabir, Mohammed J., CD286
Kaffeine media player, 784–786
Karbon14, 739–742
Kasbar taskbar, defined, 242
KAsteroids, 832
Kate
 basics, 600–602
 Kate Session Applet (KDE 4), 295
 Kate Session Menu, 233, 238
 vs. KWrite, 599
KAtomic, 839
KBackgammon, 835
KBattleship, 839
KBlackBox, 835–836
KBlocks, 832
KBounce, 833
KBovo, 836
KBreakout, 833
KChart, 743
kcmpureftpd package, CD45
KDE 3
 Adept Installer and, 929–930
 Adept Manager and. *See* Adept Manager on KDE 3
 Adept Updater and, 980
 printers, adding on (Kubuntu), 1002–1010
 printers, configuring on (Kubuntu), 1010–1012
 scanners, using on (Ubuntu), 1023–1027
KDE 3 (K Desktop Environment)
 application windows, 208–212
 backgrounds, changing, 247–248

configuration changes, undoing, 252
configuring sound on (Kubuntu), 760–762
context-sensitive menus, 215–217
customizing, 223
customizing digital device recognition on, 868–871
digital device default actions with, 866–867
Dolphin file manager. *See* Dolphin
features overview, 206–208
history of, 204–205
installing on Ubuntu, 271
K Menu, 213–214
KDE-related packages, 200
keyboard shortcuts, 252–253
KNetworkManager, 1069–1070
Konqueror browser. *See* Konqueror browser
managing users and groups on, CD116–120
menus, customizing, 217–223
mouse behavior, customizing, 224–227
mouse usage, 205
overview, 203–204
panel menus, 212–215
panels, adding new, 241–242
panels, configuring, 230–237
panels, customizing contents of, 237–241
power management, configuring, 244–246
remote printing on, CD344–345
screen resolution, configuring, 227–229
screensaver, configuring, 242–243
Strigi desktop search application, 264–267
themes, switching, 248–251
using Compiz window manager with, 200–201
window manager, changing, 268–270
wireless security on, 1072
KDE 4
Adept Installer and, 929–930
Adept Manager and. *See* Adept Manager on KDE 4
Adept Updater and, 980, 987–988
existing printers, configuring on, 1016–1017
printers, adding on (Kubuntu), 1012–1016
KDE 4 (K Desktop Environment)
application-specific menus, 280
compositing and KWin. *See* compositing
configuring sound on (Kubuntu), 763–765
context-sensitive menus, 280–284
customizing desktop, 308
customizing digital device recognition on, 871–873
digital device default actions with, 868
introduction to, 273
K Menu, 284–288
KNetworkManager, 1069–1070
KRunner, running commands with, 288–289
mouse usage, 277–278
new features in, 274–276
panel, 289–290
plasmoids (desktop widgets). *See* widgets, desktop (KDE 4)
remote printing on, CD341–344
screen resolution, configuring, 309–311
themes, switching and creating, 311
window decorations, 278–280
KDE desktop, connecting to FTP servers from, CD35–40
KDE DVI file previewer (kdvi), 611–612
KDE KAddressBook application (Contacts), 387
KDE Konsole, 546–548
KDE Remote Desktop Connection (krdc) for terminal server connections, CD25–27
KDE text editors, 597–602
kdegames package, 832
kde-hal-device-manager package, 769
kdesvn client, CD101
KDevelop IDE, CD81
kde-window-decorator, 197
KDiamond, 841
kdvi (KDE DVI file previewer), 611–612
KEdit, 597–599
Kenolaba, 836
Kerberos authentication
client, setting up, CD147–148
installing client and server commands, CD142–143
overview, CD141
server, setting up, CD145–147

kernel
 defined, 5
 kernel-level virtualization, CD206, CD208
 loading, 909
 modules, building (VMware server), CD240
 server options, configuring (NFS), CD362–363
Kernel-based Virtual Machines (KVM). *See* KVM (Kernel-based Virtual Machines)
kernel-nfs-server package, CD368
key bindings
 customizing, CD88–90
 emacs, 593
keyboard, specifying at installation, 71
keyboard commands
 emacs, 591
 gedit, 595
 KDE text editors, 597
keyboard shortcuts
 GIMP, 721
 GNOME, 156–157
 KDE 3, 252–253
 supported by KWin desktop effects, 299–301
keyword search option (-k), 559
keywords
 in boot process, 908–909
 searching for Firefox add-ons and, 445–446
 upstart job files and, 914
KFoulEggs, 833
KFourInLine, 836
KFTPgrabber, CD43
KFTPgrabber FTP client, CD43
KGoldrunner, 833
Kicker/Kicker panel
 configuration and, 230–231, 237
 first-time booting and, 47
 icons, adding and, 241
 KDE and, 200, 206, 229, 274
 panels, adding and, 241
 vs. Taskbar, 192
kids, games for (Kubuntu), 838–839
Kile, LaTeX and, 611–612, 614
killer apps, defined, CD277, CD287
KInfoCenter application, 768–769
kinit command, CD148
Kiriki, 836
Kivio, 743
KJumpingCube, 839
Klickety, 840
klist commands, CD148
Klotski, 828
KMahjongg, 836–837
KMail, configuring, 373
KMines, 840
KNetwalk, 840
KNetworkManager (Kubuntu), 1067–1070
Knuth, Donald, 606, 615
Koetter, Patrick, CD305
Kolf, 833–834
Kollision, 834
Kolor Lines, 840
Konq Profiles Applet (KDE 4), 295
Konqueror
 connecting to FTP servers in, CD36–38
 connecting to Windows shares in, CD50–52
 importing bookmarks into, 487–489
 multiple operating systems and, 475
Konqueror browser
 configuring, 261–262
 overview, 259–261
 using as file manager, 262–264
Konquest, 840
Konsole, 546–548
Konsole Profiles Applet (KDE 4), 295
Kontact
 calendar, configuring, 390–395
 calendar, managing, 416–419
 checking/reading e-mail, 396–398
 configuring, 372–373
 configuring e-mail in. *See* e-mail, configuring in Kontact
 contacts, configuring, 386–390
 contacts, creating and managing, 414–416
 filters for incoming mail, 409–411
 folders, 408–409
 importing mail into, 469–472
 launching, 371–372

PIM and, 371
preferences for reading, setting. *See* Kontact preferences, setting
spam/virus detection, configuring, 411–414
subapplications and, 371
Summary View, configuring, 395–396
switching to, 463
vs. Thunderbird, 461
Kontact preferences, setting, 398–404
e-mail, sending, 404–408
e-mail appearance, 398–401
HTML and, 401–404
threading, 401
Kooka, 1018, 1023
Kopete IM client
adding buddies in, 518–519
configuring IM accounts in, 510–513
connecting to IM accounts in, 514–517
customizing incoming message events, 522–523
exiting, 524–525
logging conversations, 523–524
organizing buddies in, 519–522
overview, 510
sources of information about, 525
using, 513–514
using IRC in, 528–530
KOrganizer. *See* Calendar within Kontact
KPhotoAlbum, 879
KPoker, 838
KPowersave applet, 244–246
kppp package, 1079
KPresenter, 703
KRandRTray application (KDE 4), 309
krdc (KDE Remote Desktop Connection) for terminal server connections, CD25–27
KReversi, 837
krfb VNC server, configuring (Kubuntu), CD9–11
Krita, 721–730
basics, 721
formats supported by, 722–723
quick tour, 724–725
resources, 729–730
starting, 723–724
using, 725–729
KRunner, running commands with, 288–289
KSameGame, 841
KSirk, 841
KSirtet, 834
KSmileTris, 834
KSnakeRace, 834
KSokoban, 841
KSpaceDuel, 834
KSpread spreadsheets, 688–695
adding KSpread, 930
basics, 689–690
Help, 690
importing spreadsheets, 693–695
invoking functions in, 690–691
starting, 688–689
type of cell data, specifying, 691–693
wizard for starting, 688–689
KSquares, 837
KSudoku, 842
Kthememanager application, 250–251
KTorrent on Kubuntu, CD61–62
KTron, 835
KTuberling (game), 838
Kubrick, 842
Kubuntu. *See also* startup process (Ubuntu/Kubuntu)
Add/Remove Programs menu item, 929–930
applications installed by default, 931–932
basic tools, 931
booting, 45–47
camera interfaces on, 867
configuring sound on, 760–772
connecting to Windows shares from, CD50–53
defined, 7
digital cameras on, 879–884
Firefox, installing on, 423
games on, 830–842
installing. *See* install options from alternate CD; installing on special-purpose systems (Ubuntu/Kubuntu); installing Ubuntu/Kubuntu

Kubuntu 8.04 and Kubuntu 8.10 Desktop CDs, 65
kubuntu-restricted-extras package, 774
migrating from, 455, 456–457
Network Manager on. *See* Network Manager on Kubuntu
overview, 14
printers, adding (KDE 3–based), 1002–1010
printers, adding (KDE 4-based), 1012–1016
printers, configuring (KDE 3-based), 1010–1012
printers, configuring (KDE 4-based), 1016–1017
readahead and, 920
recognition of digital devices on, 865–873
selecting version of, 28–29
services, configuring at startup, 919
shutdown process, 921–922
spreadsheets, 659–660
text editors. *See* Kate; KEdit; KWrite
kubuntu-desktop package, 271
kubuntu-restricted-essentials package, 786
kvim, 582
KVM (Kernel-based Virtual Machines)
creating virtual hard disk for, CD254–255
installing, CD253–254
networking for, CD257–260
overview, CD252–253
starting, CD257
using with virtual bridges, CD259–260
KWin, using compositing in, 297–301
KWin4, 837
KWord
basics, 633–634
documents, creating and editing in, 637–638
documents, importing/exporting/saving, 645
frames, 638–644
installing, 634
quick tour, 634–637
KWrite, 599–600

L

-L option, dpkg, 947
LAMP (Linux, Apache, MySQL, and Perl) servers, CD280
Lamport, Leslie, 615
languages. *See also* markup languages
support metapackages, 618
languages (spoken)
Calc and, 670–671
configuring at installation, 69–70
Impress and, 695, 696
language support metapackages, 618
LaTeX
basics, 608–610
DVI files and, 610–611
graphical tools and, 611–614
installing, 606–607
resources, 615
LaTeX: A Document Preparation System, 2nd Edition (Addison-Wesley Professional, 1994), 615
LaTeX: Line by Line: Tips and Techniques for Document Processing, 2nd Edition (1999), 615
launching. *See also* Application Launcher (KDE 4)
Evolution, 314–315
Kontact, 371–372
Launchpad, Canonical, 15
layers, defined, 707
layout, modifying, 625–627
LDAP protocol. *See* Lightweight Directory Access (LDAP) protocol
LEAP (Lightweight Extensible Authentication Protocol), 1060
lease periods (DHCP), CD314
leechers, defined (BitTorrent), CD62
Letter Wizard, 622–624
libapache2-svn package, CD101
libdbx project, 473
libraries
Library, Firefox bookmarks, 437–439
Open Clip Art Library, 706
libxine1-all-plugins package, 786
Licq IM application, 531
Liddington, Simon, CD93
Lieutenant Skat, 838
Light Craft's LightZone, 732
Lightweight Directory Access (LDAP) protocol
accessing contact information via, 352–353, 388
servers, defining, 388–389

Lightweight Extensible Authentication Protocol (LEAP), 1060
LiLo (Linux Loader), 906, 907
Line Printer Daemon (LPD), CD339–340
LinkSys WPC54G wireless PCMCIA cards, 1107
Linux. *See also* Kubuntu; Ubuntu
 accessing NFS from, CD53–54
 ACLs. *See* Access Control Lists (ACLs), Linux
 arcade games for, 851
 authentication process and PAMS, CD125–132
 backup software for, CD177–182
 command-line tools and, 535
 commands for filesystems, 552
 desktops overview, 123–125
 directories, 103–106
 emulation software for, CD208–209
 equivalents to Windows applications, 491
 files, 103–106
 files, non-Linux and, 490
 filesystems. *See* filesystems (Linux)
 firewalling, 1084–1085
 fun levels and, 912
 graphical environments for, 123–125
 IDEs for, CD81–82
 Linux distribution, defined, 5–6
 Linux for You magazine, CD108
 Linux/UNIX backup utilities, CD179
 migrating e-mail to. *See* e-mail, migrating to Linux
 net-work oriented backup for, CD180–182
 packet filtering, 1084–1085
 partitions, 73–75. *See also* partitions
 permissions, 117–121
 popular MTAs, CD288–290
 printing history, CD339–340
 reasons for using, 4–5
 RPM-based distributions, 926
 shells for, 543–544
 text editors, 570–571, 602–603
 Time Machine, CD198
 Upstart and, 917
 vector graphics tools for, 735–736
 wildcards, 563
Linux Loader (LiLo), 906, 907
Linux multimedia stack
 GStreamer framework and plug-ins, configuring, 775–776
 multimedia frameworks, 773–774
 overview, 772
 sound APIs, 773
 sound servers, 773
Linux Network Manager, 1094–1095
Linux Standard Base project, 106
Linux Terminal Server, CD28
Linux Terminal Server Project (LTSP) servers, installing, 99
Linuxant, 1108
Lisp, 594–595
Liu, Cricket, CD322
live bookmarks (Firefox), 428, 439–441
Live CDs, 27–28
loadable kernel modules, defined, 993
local backup and restore (Linux), CD179–180
local filesystems
 defined, 109
 vs. network filesystems, 110–111
 standard and journaling, 109–110
Local mailbox, 375
local storage devices, CD180
localization
 Calc and, 670–671
 Impress and, 695
 and translation, 8
log files
 /var/log/messages, CD319
 Apache, CD285–286
 CUPS, CD351
logging conversations
 in Kopete, 523–524
 in Pidgin IM client, 508–509
logical volumes (Linux), 107–109
login
 default login names, creating, 75, 76
 PAMS and (example), CD127–130
Lone Star's LONE-TAR software, CD178
lossless formats, defined, 707
lossy formats, defined, 707

LPD (Line Printer Daemon), CD339–340
LPD/LPR Host or Printer connections, 998, 1013
LPT #1 printer connection, 998, 1012
ls (list) command, 119, 537–539
ls -help command, dashes and, 540
lspci command, 1044
LTSP servers, installing, 99
Lyx, LaTeX and, 611, 613

M

MAC (Mandatory Access Control) systems, CD152. *See also* AppArmor; SELinux
MAC (Media Access Control) filtering, 1093
Mac OS X
 Adium X for, 496
 files, and Linux, 490
 printing, integrating with Ubuntu, CD347–350
 Time Machine, CD198
 VNC servers/clients for, CD12
MacMillan, Karl, CD170
Mahjongg, 828
Mahkovec, Ziga, 916
Mail Delivery Agents (MDAs), 316, CD288
mail servers, CD287–306
 basics, CD288–291
 e-mail basics, CD287
 interacting with, 321–323
 parameters, setting, 320–322
 Postfix. *See* Postfix
 Qpopper, CD298–299
 resources, CD305–306
 running your own, CD290–291
Mail Transfer Agents (MTAs)
 backups and, CD188
 basics, CD288
 parameters, specifying, 319
 receiving vs. retrieving e-mail and, 316
Mail User Agents (MUAs), 316, CD288
mailboxes
 existing and imported (Evolution), 466–467
 imported (Thunderbird), 465
Maildir mailbox, 375
mailing lists
 defining, 357–359
 Ubuntu, 22
MailMan, 357
MailScanner
 basics, CD291, CD300
 integrating with Postfix, CD303–305
mailto URLs, defined, 315
mailx, CD291
make application (UNIX), CD68–70
MAME (Multi-Arcade Machine Emulator), CD208
man command
 basics, 559
 for command help, 541–542
man pages, defined, 541
managed wireless networks, 1092
Mandatory Access Control (MAC) systems, CD152. *See also* AppArmor; SELinux
Manifesto (Ubuntu), 7–8
manually migrating e-mail, 473–475
Marketplace (Ubuntu), 24
markup languages. *See also* HyperText Markup Language (HTML); LaTeX; TeX
 document, 606
marshalling, defined, CD357
Mashey, John, 543
masters, defined, 1027
matching e-mail messages, filtering and, 339–341
Mattessich, Richard, 660
Mayer, Frank, CD170
mbox
 basics, 458–459
 converting directories to, 474
McCarty, Bill, CD170
McCool, Rob, CD280
McLoughlin, Mark, CD7
McNealy, Scott, 615
MDAs (Mail Delivery Agents), 316, CD288
media
 backing up to removable, CD182
 for backups, CD173
Media Access Control (MAC) filtering, 1093
Meebo IM application, 531
meetings, creating and editing (Calendar), 364–365
menus. *See also* context-sensitive menus
 boot loaders, 907, 908
 Calc, 672, 673
 Draw, 733
 games (Kubuntu), 841–842
 Gnumeric, 681
 Impress, 699
 Inkscape, 737
 Karbon14, 740
 Krita, 724
 KSpread, 689
menus (GNOME)
 context-sensitive, 132–134
 customizing, 134–138
 panel, 130–132
menus (KDE 3)
 context-sensitive, 215–217
 customizing, 217–223
 panel, 212–215

metadata, defined, 708
Metal Blob Solid arcade game, 855
Meta/Super label key, 301
Metcalf, Bob, 1050
MicroEmacs, 570
MicroLITE's BackupEDGE and RecoverEDGE software, CD178
Microsoft Exchange servers, 314
Microsoft Outlook. *See* Outlook
Microsoft Outlook Express. *See* Outlook Express
Microsoft Windows. *See* Windows
MID (Ubuntu Mobile Internet Device) Edition, 16
migration. *See also* Windows systems, migrating from
 backing up personal data for, 457–458
 checking system before, 477–478
 dual-boot systems and, 489–491
 manual, 459
 profile migration, troubleshooting, 478–479
 via Thunderbird. *See* Thunderbird, migrating e-mail via
 Windows application data, 489–491
Mines, 829
Minesweeper (Windows game), 829
MITRE Polgen Policy Generation Tools, CD169
MITRE SLAT (SELinux Analysis Tools), CD169
mkfs command, 1039
mkswap command, 1039, 1043
Mobily, Tony, CD286
modal editors
 defined, 573
 vs. modeless text editors, 571
modeless editors, defined, 571
modems and dial-up connections, configuring, 1079
modes (emacs), 594
Monitor & Display configuration application (KDE), 228–229
monochrome, defined, 708
monopd server, 850
Monopoly clones (games), 847–851
Mosaic Web browser, 421–422
Motor IDE, CD81
mount options, 113, 116
mount points
 Linux filesystems and, 106–107
 partitions and, 88, 89–90
mounting
 defined, 107
 Linux filesystems, 111–114
 Windows shares at boot time, CD47–48
mouse
 behavior, customizing (GNOME), 138–140
 behavior, customizing (KDE 3), 224–227
 using (Firefox), 426
 using (GNOME), 126
 using (KDE 3), 205
 using (KDE 4), 277–278
moving
 directories, 257–258, 552–556
 files, 257–258, 552–556
Mozilla Foundation, 422
Mozilla's Thunderbird webmail extension, 348
MP3 files, 490–491
MPG files, 490
mplayer application, files and, 490–491
MPlayer Movie Player, 774
MTAs. *See* Mail Transfer Agents (MTAs)
MUAs (Mail User Agents), 316, CD288
multi-homed systems, 1053
multimedia frameworks (Linux), 773–774
multimedia stack (Linux). *See* Linux multimedia stack
multiple audio output devices, 763, 765
multi-protocol IM clients, 496
music
 adding to iPods, 897–898
 managing with Amarok, 888–896
 managing with Rhythmbox, 885–888
Muslim Edition, Ubuntu, 14
mv command, 554–555, 556
mydestination parameter, CD298
myorigin parameter, CD298
MySpace Messenger, 500
MySpaceIM, 500

N

n (new) command, 1040
names
 default login names, creating, 75, 76
 hard drives, 1029
 identifying commands by, 561
 SATA, problems with name changes, 1031
 for server, specifying, 71–72
naming
 e-mail accounts (POP), 324
 printers, 1000–1001, 1008–1009, 1015, 1016
nano, 602, 603
NAS (Network Audio System), 773
NAT (Network Address Translation), networking (virtual machines), CD247
nat, VirtualBox and, CD227
Nautilus file manager (Gnome)
 accessing existing partitions with, 50–52
 basics, 157–162
 burning data CDs/DVDs with, 799–801
navigating
 in emacs, 591–592
 in vi, 574–575
navigation toolbar (Firefox), 424
NCSA httpd (HTTP Daemon) Web server, CD280
ndisgtk package, 1105
ndiswrapper, 1104–1105, 1106, 1107–1108
NetBeans IDE, CD82
NetBIOS, CD372
netfilter framework, 1084–1085
netmask, defined, 1052
Netscape Navigator, 422
Network Address Translation (NAT)
 networking (virtual machines), CD247
 VirtualBox and, CD227
Network Audio System (NAS), 773
network filesystems (Linux), 110–111
Network Folder Wizard (KDE), CD37, CD51
Network Information System (NIS), CD360
network interfaces, Xen configuration files for, CD265
Network Manager, 1094–1095
Network Manager on Kubuntu
 configuring wired/wireless connections on KDE 3, 1071–1074
 configuring wired/wireless connections on KDE 4, 1074–1076
 KNetworkManager overview, 1067–1070
Network Manager on Ubuntu
 configuring VPN connections, 1062–1067
 configuring wired connections, 1056–1057
 configuring wireless connections, 1057–1061
 overview, 1054–1055
net-work oriented backup for Linux, CD180–182
Network Time Protocol (NTP), configuring, CD143–145
networking
 KVM support for, CD257–260
 VirtualBox and, CD227–232
network-manager-openvpn package, 1065
networks
 basics, 1050–1053
 configuration files for interfaces, 1077–1079
 configuring hardware, 1053–1054
 copying files over, 56
 firewalls, installing and configuring, 1083–1090
 Linux firewalling and packet filtering, 1084–1085
 Network Manager on Kubuntu. *See* Network Manager on Kubuntu
 Network Manager on Ubuntu. *See* Network Manager on Ubuntu
 overview, 1049–1050
 system security tips, 1082–1083
 testing with Network Tools application, 1080–1082
 virtual machines and, CD246–248
NeXT, CD280
NFS (Network File System)
 basics, 111
 configuring common options for, CD361–362
 configuring kernel server options, CD362–363
 debugging problems, CD367
 directories, accessing from Linux, CD53–54
 directories, mounting on clients, CD365–366

exported directories, defining in /etc/exports, CD363–365
exported filesystems, verifying, CD368–369
home directories, automounting, CD367
information references, CD370
operations, debugging, CD369–370
overview, CD355–357
pseudo-filesystem and exported directories, CD363
RPC services, verifying, CD367–368
servers, configuring, CD360
servers, installing, CD360
users and authentication, CD360
versions of, CD359
workings of, CD357–359
Nibbles, 830
NIS (Network Information System), CD360
nmap application, 1083
Noatun media player, 784
NoMachine NX Server, CD27
nonstandard repositories, 942
Notes (KDE 4), 296
Notification Area/System Tray applet (KDE 3), 208
Now Playing (KDE 4), 296
NTP (Network Time Protocol), configuring, CD143–145
ntpdate application, CD144–145
NX server, CD28

O

OEM mode, installing in, 97–98
office-related software. *See* Calc; Impress; KPresenter
Oikarinen, Jarkko, 494
online calendars, integrating, 366–367
online help (Evolution), 329
online photo services, exporting to, 877–878
Open Clip Art Library, 706
Open Office suite. *See also* OpenOffice.org
basics, 605, 615–616
languages for, 617–618
versions of, 616
open source software, CD70–71
OpenBSD project, CD5
OpenGL (Open Graphics Library), 175, 176–177
OpenGL plug-in, CD76
OpenOffice.org. *See also* Calc
documents, 617
tools, 491
OpenOffice.org presentations. *See also* Impress
creating, 695–703
KPresenter, 703
OpenOffice.org spreadsheets, 669–680
basics, 669–670
Calc. *See* Calc
files, installing, 670–671
Gnumeric spreadsheets. *See* Gnumeric spreadsheets
KSpread spreadsheets. *See* KSpread spreadsheets
OpenOffice.org Writer
basics, 615–616
documents, importing, 630–633
file extensions, 617
installing files for, 617–618
personalizing, 620–622
quick tour, 618–620
styles and layout, modifying, 625–627
templates, 627–630
wizards for document creation, 622–625
OpenSSH server, CD5
OpenSSL library, 1065
OpenVPN, setting up, 1065–1067
operating systems. *See also* JEOS (Just Enough Operating System)
installing in virtual machines, CD220–223, CD255–256
multiple operating systems and Konqueror, 475
Portable Operating System Interface (POSIX) standard, CD135
options
Alternate Install CDs boot menu, 95, 96
defined, 537
Server Install CD, 67
Orca application, 155
OSS sound API (Linux), 773

OSXvnc server, CD12
Outlook
 defined, 458
 migration via Thunderbird. *See* Thunderbird, migrating e-mail via
Outlook Express
 defined, 458
 e-mail, migrating with readdbx, 473–474
 e-mail, migrating with readpst, 474–475
 manual migration, 459
 migration via Thunderbird. *See* Thunderbird, migrating e-mail via
output redirection, 563–564

P

p (print) command, 1040, 1041
P2P file sharing (Ubuntu)
 BitTorrent, CD58–63
 gtk-gnutella, installing/using, CD56–58
 overview, CD55–56
package management software, 931–932
packages. *See also specific* packages; system software packages
 adding for tools support, 67
 converting from other formats, 988–991
 init packages, 911
 language support metapackages, 618
 source code, retrieving with apt-get, 952–953
 Synaptic Package Manager. *See* Synaptic Package Manager
packet filtering (Linux), 1084–1085
packets, defined, 1051
padevchooser package, 776–777
Pager (KDE 4), 289, 296
Pager applet (KDE 3), 208
paid support, Canonical Ltd. (Ubuntu), 23–24
palettes
 Inkscape, 738
 Karbon14, 741
 Krita, 725
PAMs (Pluggable Authentication Modules)
 configuration files, CD130–131
 configuration files for applications/services, CD126–127
 default configuration file, CD131–132
 overview, CD125
 PAMS used by login process (example), CD127–130
panel menus
 GNOME, 130–132, 142–147
 KDE 3, 212–215
panels (KDE 3)
 adding new, 241–242
 configuring, 230–237
 customizing contents of, 237–241
 hiding automatically, 232
panels (KDE 4), 289–290
pantone colors, defined, 708
Parallel Realities, 855
parallel virtual machines, CD205–206
parameters
 MTA, 319
 setting for mail servers, 320–322
paravirtualization, CD206
paravirtualized systems
 domainU system (Xen), CD267–272
 GUIs in, CD270
paravirtualized Xen virtual machine configuration files, CD264
parser errors (DNS servers), CD336
parted tool, 1033
partitioning
 disks, manually, 42–45
 hard drives, 1032–1033, 1039–1043
 layout, manually specifying, 84–93
 reasons for (Linux), 107
partitions
 accessing from CD, 50–56
 command-line tools to identify, 1038–1039
 creating, 1035–1037
 GNOME Partition Editor and, 133–137
 Linux and, 1035
 Linux filesystems and, 106–107
 size of, and logical volumes, 107–108

partitions (Linux), 111
 configuring manually, 73–75
 disk drives, 107
 mounting at boot time, 114–117
 mounting filesystems and, 111–114
partner repository component, 934
passwords
 backing up, 477
 BackupPC package, CD188–189
 Firefox and, 433–434
 user passwords, creating, 77–78
PATA (Parallel Advanced Technology Attachment) drives, 1027, 1029
PATA/SATA, problems after adding to SCSI systems, 1032
Patel, Neil, 192
PATH environment variable, 567
pathnames, defined, CD278
Patience, 838
PCI cards, adding, 1044–1045
PCL, defined, 708
PCMCIA cards, 1043–1044
PDD files for printers, 999, 1000, 1001
PDF files, 490–491
Peer-to-Peer Protocol (PPP), 1079
PenDriveLinux.com, 57
permissions
 AppArmor, CD159
 chmod command and, 556
 directory (Linux), 119–120
 files and directories, changing, 556–558
 Linux, 117–121
 umasks and, 558
Persistence function (Live CD), 27, 56–60
Personal Computer Memory Card International Association. *See* PCMCIA cards
personal data, backing up for migration, 457–458
Personal Information Management (PIM) package
 Evolution support for, 313
 Kontact and, 371
Personal Storage Table (PST) format
 files, migration and, 457
 migrating with readpst and, 474–475
Phonon (KDE 4), 276
Phonon Multimedia API, 773
photo albums, 880
photos
 editing (F-Spot), 876–877
 editing (Krita), 730–733
Photoshop Express, 732
physical volumes, defined, 108
Picasa photo management software, 730, 732
pico, 602
Picture Frame (KDE 4), 296
Pida IDE, CD82
Pidgin IM client
 adding buddies in, 505–507
 configuring IM accounts in, 498–501
 connecting to IM accounts in, 504–505
 exiting, 509
 logging conversations, 508–509
 online help for, 509
 organizing buddies in, 507
 overview, 498
 using, 501–504
 using IRC in, 526–528
 viewing buddies in, 507–508
 Windows version, 496
PIM (Personal Information Management) package
 Evolution support for, 313
 Kontact and, 371
pipeline, defined, 751
pipes, 563–564
pixels, defined, 708
pkg-config package, CD74
PKI (public key infrastructure), CD5
Place Windows plug-in, 191
Places menu (GNOME), 131
Planet Ubuntu blog, 20
playlists
 creating with Amarok, 816–817
 creating with Rhythmbox, 810–811
 creating/editing on iPods, 898–900
 managing with Amarok, 894–895
 managing with Rhythmbox, 887–888

Pluggable Authentication Modules. *See* PAMs (Pluggable Authentication Modules)
plug-ins
 AWN, 194
 Compiz, 178
 Compiz window manager, 192
 Subversion for Eclipse, CD99
plug-ins (Firefox)
 adding, 450–453
 detection of missing, 427–428
PNG (Portable Network Graphics), 708
Podcasts
 adding with Amarok, 819
 adding/playing with Rhythmbox, 813–814
Point-to-Point Tunneling Protocol (PPTP), 1062
POP (Post Office Protocol). *See also* POP3
 accounts, naming, 324
 Gmail, accessing with, 345
 Gmail, configuring for, 343–345, 346
 vs. IMAP, CD294, 318–322, 374
 mail, sending and, 323–324, 347
 mail servers, interacting with, 321–323
 MDAs and MUAs and, CD288
 migrating e-mail and, 459
 POP mailers settings, 320–321
 time zone, 325–326
 Yahoo!, access to, 349
POP3
 accounts for receiving e-mail, 377–379
 defined, 374
 vs. IMAP, 374
 migrating e-mail and, 459
pop3check application, CD93
popd command, 552
pop-ups, Firefox, 428, 430–431
portable audio devices, configuring in Amarok, 888–889
portable hard drives, 1027
Portable Operating System Interface (POSIX) standard, CD135
portmapper daemon, CD358
ports
 remapping for VirtualBox VM access, CD232
 SATA and, 1029
Post Office Protocol. *See* POP (Post Office Protocol); POP3
Postfix
 adding spam filtering to, CD300–305
 adding virus scanning to, CD300–305
 basics, CD289
 configuration models, CD292–294
 configuring, CD294–298
 installing, CD291–294
 postfix-doc, CD292
 related packages, CD291–294
Postfix (Sams, 2001), CD305
Postfix: The Definitive Guide (O'Reilly Media, 2003), CD305
postgrey, CD292, CD300–301
PostScript, 708
Potato Guy, 838–839
power management, configuring
 GNOME, 148–152
 KDE 3, 244–246
PPP (Peer-to-Peer Protocol), 1079
PPT files, working with, 491
PPTP (Point-to-Point Tunneling Protocol), 1062
preferences. *See also* Kontact preferences, setting
 configuring (Firefox), 449–450
 configuring (gedit), 596–597
 configuring (Synaptic Package Manager), 961–962
 e-mail, sending and receiving, 331–333
 Software Preferences tool, 940
Preferences (Firefox), 429, 430
preformatted print jobs, CD352
prepress, defined, 708
presentations. *See also* Impress
 KPresenter, 703
previews, scanners, 1020, 1021, 1023, 1024–1025
print quotas, defining, 1008
print servers
 CUPS, CD340–345, CD351–352
 Linux/Unix printing history, CD339–340
 Mac OS X printing, integrating with Ubuntu, CD347–350
 preformatted print jobs, CD352
 remote printing on KDE 3, CD344–345

remote printing on KDE 4, CD341–344, CD350–352
Windows printing, integrating with Ubuntu, CD346–347
printenv command, 567
Printer Configuration application, 1016
Printer Wizard, 1002–1010
printers
selecting models, 999–1000, 1014–1015, 1016
sharing using Samba, CD376–377
printers, adding
basic configuration, 995
basics, 994–995, 1002
configuring existing (KDE-based Kubuntu), 1010–1012, 1016–1017
configuring existing (Ubuntu), 995–997
on KDE 3-based Kubuntu, 1002–1010
on KDE 4-based Kubuntu, 1012–1016
manually defining new (Ubuntu), 997–1002
Printer Wizard, 1002–1010
privacy settings, configuring (Firefox), 431–434
privileged operations (Ubuntu), 120–121
Pro DNS and BIND (Apress), CD322
processors supporting virtualization, CD207
profiles
AppArmor, CD153, CD157–161
migrating existing (Firefox), 476–478
prompts, defined, 535, 536
properties, customizing panel (GNOME), 143
protocols, Web servers and, CD278
pseudo-filesystem (NFS), CD363
Psi IM application, 531
PST (Personal Storage Table) format
files, migration and, 457
migrating with readpst and, 474–475
public key infrastructure (PKI), CD5
publishing, free/busy information (calendars), 367–369
publishing documents. *See* documents, creating and publishing
PulseAudio sound server, configuring, 776–779
pure-ftpd package, CD45
pushd command, 552
PVM clustering software, CD205–206
pwd command (process working directory), 552

Q

QEMU CPU emulator
installing, CD271–273
KVM virtual machine and, CD254–255
user networking and, CD258
Qmail, CD289
Qpopper, CD292, CD298–299
QQ software, 500
Qt software, 204
qtparted tool, 1033
Quackle (Scrabble clone), 851
query/replace operations (vi), 578
question mark (?) wildcard, 562
Quick browser menu (KDE 3), 234
Quick Launcher, 206, 239, 259, 372, 423
quick tours
Draw, 733–735
GIMP, 711–713
Inkscape, 737–738
Karbon14, 740–741
Krita, 724–725
Quicklauncher applet (KDE 3), 208
QuickStart menu (KDE 3), 234

R

Radack, Steven, CD93
radio stations. *See* Internet radio stations
RADIUS (Remote Authentication Dial-In User Service), 1060
RAID (Redundant Array of Inexpensive Disks), vs. backups, CD172
rapidsvn client, CD101
raster graphics
defined, 708
formats, 735–736
RBAC. *See* Role-based Access Control (RBAC)
rdesktop for terminal server connections, CD20–22
RDP (Remote Desktop Protocol), CD20, CD212
readahead, 920

readdbx, migrating e-mail with, 473–474
readline library, defined, 560
readoe utility, 473
readpst, migrating e-mail with, 474–475
Really Simple Syndication (RSS), live bookmarks and, 439
RealVNC, CD7
RecoverEDGE software, CD178
redirection, input/output, 563–564
reducing images (GIMP), 718–719
Redundant Array of Inexpensive Disks (RAID), vs. backups, CD172
reformatting USB Flash drives, 59
Refresh rate (screens), 227–228
Reiser, Hans, 110
Remote Authentication Dial-In User Service (RADIUS), 1060
remote connections software, CD27
Remote Desktop Connection application (Kubuntu), CD16–19
Remote Desktop Protocol (RDP), CD20, CD212
Remote Desktop Viewer application (Ubuntu), CD13–16
remote printing
 on KDE 3, CD344–345
 on KDE 4, CD341–344
 troubleshooting, CD350–352
Remote Procedure Calls (RPCs)
 basics, CD357–358
 services, verifying (NFS), CD367–368
remote systems, logging into, CD4–5
removable media filesystems, mounting, 117
removing software. *See* software, adding/removing/updating
renaming, files and directories, 161, 258, 552, 554–555
rendering, defined, CD278
replacing
 text, in emacs, 592
 text, in vi, 577–578
 x86, 592
replying to e-mail, 398
repositories and components, 932–942
 additional components, enabling, 934–936
 alternative repositories, 936–942
 backports repository component, 934
 components, 933–934
 nonstandard repositories' problems, 942
 partner repository component, 934
 repositories, defined, 932
 repositories, organizing, 933
 Software Sources tool, 936–942
 Subversion repository. *See* Subversion
rescue mode, booting and, 68
Resize option, for disks, 74
resolution
 defined, 708
 scanners, 1020–1021, 1024
resolvers, CD322, CD323
resources. *See also* Web sites for further information
 Apache, CD286
 emacs, 595
 mail servers, CD305–306
 SELinux, CD170
 TeX and LaTeX, 615
 vi and vim, 583
 Web servers, CD286
restarting DNS servers, CD333–335
restoring
 from backups, CD195–198
 files, CD182–184
 local restore, CD179–180
retrieving e-mail
 receiving vs. retrieving, 316, CD288
 security, defining for, 380–381
reverse lookup files
 creating, CD331–332
 incorporating with BIND, CD332–333
reverse map files. *See* reverse lookup files
RFCs, CD321–322
RGB, defined, 709
Rhythmbox
 creating playlists, 810–811
 digital audio players and, 885–888
 playing audio files/Internet radio, 809–814
 playing CDs, 780

playing/adding Internet radio stations, 811–813
playing/adding Podcasts, 813–814
ripping CDs with, 787–789
Ring Switcher plug-in (Compiz), 192
ripping CDs
with k3b, 789–792
with Rhythmbox, 787–789
rm command, 555, 556
roaming, defined, 1092
Robots, 830
Rõckai, Peter, 972
Role-Based Access Control (RBAC)
basics, CD152
emulating with AppArmor, CD166–167
root directories, defined, 104
root kits, 1083
Rosetta translation system, 8
Rotate Cube plug-in (Compiz), 188–192
routers, 1051–1052
rows, headings (Calc), 672
RPCs (Remote Procedure Calls)
basics, CD357–358
services, verifying (NFS), CD367–368
RSS (Really Simple Syndication)
live bookmarks and, 439
RSS feed, defined, 439
rsync
backing up with, CD185–187
as backup transport, CD187
basics, CD182
rsyncd
as backup transport, CD187
defining backups with, CD191–193
rulers
Draw, 734
Inkscape, 738
rules
AppArmor, CD153, CD157–161
SELinux, CD153
SpamAssassin, CD305
run levels, init process, 912

S

-S option (dlocate utility), 949
Safari
bookmarks and, 484–485
migrating from, to Firefox, 479–481
Safari Bookmark Exporter, 484–485
Samba (Server Message Block) software, CD45–46
Samba server
configuring, CD374–379
information references, CD380
installing, CD373–374
overview, CD373
Samba 3 By Example, CD374
sharing directories via, CD379–380
sharing printers/home directories using, CD376–377
Same Gnome, 830
SamePlace IM application, 531
Sametime software, 500
SANE (Scanner Access Now Easy)
basics, 1018
sane-findscanner, 1022
sane-utils package, 1022
SATA (Serial Advanced Technology Attachment) drives
adding, 1029
vs. PATA, 1029
problems after adding IDE drives to, 1031
problems after adding to SCSI systems, 1032
problems with name changes, 1031
Satellite system (Postfix configuration model), CD293
Sauerbraten arcade game, 857
Sbcl IDE, CD82
scaling images (GIMP), 718, 720
scanners. *See also* SANE (Scanner Access Now Easy)
using on KED 3-based Ubuntu, 1023–1027
using on Ubuntu, 1018–1022
ScatterChat IM application, 531–532
SCCS (source code control system), CD99–100
scheduling Calendar events, 360–361, 417–418
schemes, defined, CD278

scopes, CD315, CD316
Scrabble clones, 851–852
screen movement commands (vi), 575
Screen Resize & Rotate application (KDE 4), 310
screen resolution, configuring
 GNOME, 140–142
 KDE 3, 227–229
 KDE 4, 309–311
screensaver, configuring
 GNOME, 147–148
 KDE 3, 242–243
screenshots (GIMP), 714–718
Scribus, 645–656
 basics, 645–646
 installing, 646–647
 resources, 656
 tutorial, 647–656
scripts
 configure scripts, CD73
 startup command scripts, 912, 913
SCSI
 drives, adding, 1030
 systems, problems after adding PATA/SATA drives to, 1032
SD (Secure Digital) memory cards, 901
SeaMonkey, 422
search folders (e-mail), 336–338
searches
 Beagle desktop search application, 162
 for emacs, 945–946, 947
 for packages, 946, 947
 for software, 962–964, 976
 for software updates, 977–979, 986–987
 Strigi desktop search application, 264–267
 Synaptic, 962–964
 text, in emacs, 592
 text, in vi, 577–578
 Tracker search application, 165–166
second-stage boot loaders, 907
Secure Digital (SD) memory cards, 901
Secure Sockets Layer/Transport layer Security (SSL/TLS), 1065
security. *See also* AppArmor; SELinux
 for composing e-mail (Kontact), 407
 defining for retrieving e-mail, 380–381
 defining for sending e-mail, 383–384
 Linux vs. Windows counterparts, CD151
 mail servers and, CD291
 security mechanisms, CD152–153
 Security pane (Firefox), 432, 433
 system security tips, 1082–1083
 wireless networking, 1093–1094
 wireless on KDE 3, 1072
Security Enhanced Linux Symposium — SELinux 2007 (DEStech), CD170
seeding, defined (BitTorrent), CD58
SEEdit SELinux Policy Editor, CD169
selection boxes, defined, 662
SELinux
 vs. AppArmor, CD152–154
 basics, CD152, CD167–168
 installing, CD168–169
 resources, CD170
 tools, CD169–170
SELinux by Example: Using Security Enhanced Linux (Prentice-Hall, 2006), CD170
SELinux: NSA's Open Source Security Enhanced Linux (O'Reilly, 2004), CD170
Sendmail
 basics, CD289–290
 Evolution support of, 323–324
sequence numbers, identifying commands by, 560
Serial Advanced Technology Attachment. *See* SATA (Serial Advanced Technology Attachment) drives
Serial Port #1 printer connection, 998, 1012
Server Edition (Ubuntu), 15–16
Server Message Block. *See* SMB (Server Message Block)
servers
 booting, 66–67, 93
 configuring (NTP), CD143–144
 configuring (VNC), CD7–12

Install CD options. *See* Ubuntu Server Install CD options
Install CD (Ubuntu), 18–19
installing (Ubuntu). *See* configuring at installation (Ubuntu)
NFS, configuring, CD360
NFS, installing, CD360
print servers. *See* print servers
setting up (Kerberos), CD145–147
setting up (Subversion), CD103–106
sound (Linux), 773
Service Set Identifier (SSID), 1059, 1093
services, configuring, 918–919
shell scripts, defined, 543
shells, 542–544
shells, starting, 544–551
with GNOME terminal application, 545–546
with Konsole, 546–548
with X Window System's terminal application, 548–551
Shift Switcher plug-in (Compiz), 179
Shisen-Sho, 837
Shorten (SHN) audio codec, 822, CD78
Shoutcast, 818
shutdown (Ubuntu/Kubuntu), 921–922
Shuttleworth, Mark, 6
sidebars (Firefox), 427
Simple Instant Messenger (SIM-IM), 500
Simple Mail Transport Protocol (SMTP)
Evolution support of, 323–324
sending outgoing Gmail, 345–347
Sitter, Harald, 872
slash (/), 537
SlashDot, 297
slaves, defined, 1027
sleep mode, 150
Smart Backup, CD178
smarthost, CD292
SMB (Server Message Block)
as backup transport, CD187
defining backups with, CD193–194
networking protocol, CD372–373
SMB/CIFS filesystems, 113–114
SMTP. *See* Simple Mail Transport Protocol (SMTP)
SOA records, zone files and, CD328–329
software. *See also* system software packages
add-ons, building, CD78–79
development software, CD66
device drivers, defined, 993
for Linux backups, CD177–182
minimizing for startup, 918–919
office-related. *See* Calc; Impress; KPresenter
open source, CD70–71
selecting server software for installation, 80, 81
source code control. *See* source code control software
standalone, building, CD71–78
supported, building, CD79–80
software, adding/removing/updating, 925–992
Add/Remove menu item (Ubuntu), 927–929
Add/Remove Programs menu item (Kubuntu), 929–930
Adept tools. *See* Adept Package Manager
with apt-get. *See* apt-get
with aptitude. *See* aptitude
basics of adding/removing, 926
exploring system software. *See* system software packages
packages, converting from other formats, 988–991
repositories and components. *See* repositories and components
software packages management, 931–932
with Synaptic. *See* Synaptic Package Manager
systems, keeping efficient, 991–992
Ubuntu community, 925
Ubuntu Update Manager, 970–972
Software Sources tool, 936–942
Solid (KDE 4), 276
sound
APIs (Linux), 773
configuring on KDE 3, 760–762
configuring on KDE 4, 763–765
configuring on Kubuntu, 760–772
configuring on Ubuntu, 750–760
sound servers (Linux), 773

system levels, setting (Kubuntu), 765–766
system levels, setting (Ubuntu), 753–755
system sound devices, specifying and testing (Ubuntu), 751–753
system sound devices, specifying (Kubuntu), 760–764
system sounds, testing/customizing (Kubuntu), 766–768
system sounds, testing/customizing (Ubuntu), 755–756
Sound Converter, installing/using, 820–822
soundKonverter, installing/using, 822–824
source code control software, CD99–100. *See also* Subversion
source code control system (SCCS), CD99–100
source code, retrieving, CD67
SourceForge site, CD20
Sox project, 803
spam
detecting (Kontact), 411–414
filtering, 338–341, CD300–305
SpamAssassin
basics, 341–342, 414, CD300
integrating with Postfix, CD303–305
special effects (Compiz), 178
specialty platforms, Ubuntu versions for, 16
spot backups, CD174
spreadsheets, 659–695
basics, 660
calculations, 661–666
Kubuntu, 659–660
OpenOffice.org. *See* OpenOffice.org spreadsheets
transferring to Ubuntu, 677
tutorial, 660–669
Ubuntu, 659–660
Spyglass Mosaic, 422
ssh application for logging into remote systems, CD4–5
SSH protocol, CD5–6
SSID (Service Set Identifier), 1059, 1093
SSL/TLS (Secure Sockets Layer/Transport layer Security), 1065
stacks, defined, 552
Stallman, Richard, 3, 570
standalone software, building, CD71–78
standard local filesystems, 109–110
StarOffice, 616
starting
Calc, 671–672
Draw, 733
emacs, 583–584, 588–591
Firefox, 422–424
GIMP, 710–711
Gnumeric, 681
Impress, 696–698
Inkscape, 736–737
Karbon14, 739–740
Krita, 723–724
KSpread, 688–689
vi, 572
startup process (Ubuntu/Kubuntu), 905–916
boot loader, 906–909
boot monitor or BIOS, 906
init process. *See* init (initialization) process
initrd or initramfs, 909–911
kernel, loading, 909
Steele, Guy, 570
stemming, defined, 162
sticks, creating bootable, 57–60
Sticky Notes applet (GNOME), 145
storage
adding, 1027
devices, 106–107, 1027
logical volumes and, 107
Thunderbird mail storage, 461
Storix's System Backup Administrator, CD178
Strigi desktop search application, 264–267
Studio (Ubuntu), 14–15
Styles and Formatting window (Impress), 699
styles of documents, modifying, 625–627
su command, vs. sudo command, 120–121
subapplications (Kontact), 371
Subversion
information references, CD107–108
installing, CD100–101
plug-ins for Eclipse, CD99

repository, checking files out of, CD106
repository, committing changes to, CD107
repository, importing projects into, CD102–103
repository, setting up, CD101–102
server, setting up, CD103–106
sudo command
apt-get command and, 950
basics, CD132–134
fdisk command and, 1039
privileged operations and, 120
recovering broken, CD134–135
vs. su command, 120–121
Sudoku, 829
superusers, defined, 120
support for Ubuntu/Kubuntu, 19–24
supported software, building, CD79–80
suspend (sleep mode), 150
SVG
SVG format (Draw), 733, 736
SVG images, importing (Krita), 725, 726
SVG standard, 735
svnserve server, CD103, CD105–106
swap partitions, 85, 91–92
SWF files, working with, 490
Synaptic Package Manager
vs. Adept Manager tools, 952
vs. apt-get, 950
basics, 931, 932, 959–961
build environment core elements and, CD66
internationalizing Calc and, 670–671
packages, deleting in, 968–970
packages, installing in, 964–968
preferences, configuring, 961–962
shells, installing and, 543
software, adding/removing and, 959–970
software, searching for, 962–964
Ubuntu and, 659
syntax
AppArmor profiles, CD158–159
DNS servers errors, CD336
System Backup Administrator, CD178
system groups. *See* groups
System menu
GNOME, 131–132
KDE 3, 214–215
System Menu applet (KDE 3), 208
System Monitor applet (GNOME), 145
System Settings application (K Menu), configuring themes with, 248
system software packages, 942–950
listing available packages, 943–946
listing content, 947–948
listing information, 947
packages providing missing files, determining, 949–950
packages providing specific files, determining, 948–949
system sound
devices, specifying (Kubuntu), 760–764
levels, setting (Kubuntu), 765–766
levels, setting (Ubuntu), 753–755
testing and customizing (Kubuntu), 766–768
testing and customizing (Ubuntu), 755–756
system startup and shutdown, 905–923
Bootchart, 916–917
optimizing boot process, 917–920
shutdown (Ubuntu/Kubuntu), 921–922
startup process (Ubuntu/Kubuntu). *See* startup process (Ubuntu/Kubuntu)
system types supporting Ubuntu/Kubuntu, 17
system-config-printer, adding printers and, 995
system-config-samba tool, CD374
system-level virtualization, CD205
systems
installing command-line systems, 98
installing desktop systems in text mode, 96–97
installing in OEM mode, 96–97
keeping efficient, 991–992
security tips, 1082–1083
upgrading with apt-get, 951–952
SysVInit system initialization
init process, 913
SysVInit package, 911, 912
traditional, 913

T

-t option, mount command and, 113
Tab key, command and filename completion and, 561
tabbed browsing (Firefox), 427
tactics and strategy games (Kubuntu), 839–841
tags
 bookmarks (Firefox), 436–437
 editing (ID3), 885
 KDE 3 Adept Manager, 974, 975
Tali, 827
tar
 for archiving and restoring files, CD180, CD182–184
 as backup transport, CD187
 tar-format archive, CD197
tarballs, CD183–184
Task Manager (KDE 4), 289, 296
Taskbar applet (KDE 3), 208, 236–237
TECO (Text Editor and Corrector), 570
telnet application, CD3–4
templates
 creating new Scribus documents and, 648–649
 OpenOffice.org and, 627–630
Temporary Key Integrity Protocol, 1094
terminal server connections
 krdc for, CD25–27
 rdesktop for, CD20–22
 tsclient for, CD22–25
Terpstra, John, CD374
testing
 backups, CD176–177
 DNS servers, CD333–335
 with Network Tools application, 1080–1082
 printers, 1010, 1015
 Samba availability and services, CD379
 sound devices, 751–753
 system sounds, 755–756, 766–768
testparm utility, CD378–379
Tetravex, 829
TeX
 basics, 606–608
 DVI files and, 610–611
 graphical tools and, 611–614
 resources, 615
TeX Live, 606
TeXbook (Addison-Wesley Professional, 1984), 615
TeXmacs package, 614
text
 deleting and changing in vi, 575–576
 inserting in vi, 574
 searching and replacing (emacs), 592
 searching and replacing (vi), 577–578
Text Editor and Corrector (TECO), 570
text editors. *See also specific* text editors; text files, editing
 DHCP configuration files and, CD311–313
 Linux, 570–571, 602–603
text fields, adding to cells, 666
text files, editing, 569–603
 basics, 569–570
 emacs. *See* emacs
 gedit, 595–597
 KDE text editors, 597–602
 Linux text editors, 570–571
 modal vs. modeless editors, 571
 various editors listed, 602–603
 vi. *See* vi
text mode
 installing desktop systems in, 96–97
 installing in OEM mode and, 97–98
Text Mode installer (Calc files), 670
Text Run Around options (KWord), 640
THE (The Hessling Editor), 603
themes
 AWN, 195
 GNOME, 153–155
 switching and creating (KDE 4), 311
themes (Firefox)
 adding, 444–450
 viewing, 442–444
themes (KDE 3)
 installing with Kthememanager, 250–251
 switching, 248–251
Thompson, Ken, 543, 570

threading, setting (Kontact), 401
Thunderbird, migrating e-mail via, 459–472
 basics, 459–461
 importing into Evolution, 466–469
 importing into Kontact, 469–472
 installing/updating Thunderbird, 462
 vs. Kontact or Evolution, 461
 mail storage, 461
 users, existing, 463–466
 users, new, 462–463
TIFF, defined, 709
TightVNC, CD7, CD12
Time Machine, Linux, CD198–201
time machines, Linux, CD198–201
Time To Live (TTL), zone files and, CD327–329
time zones
 configuring at installation (Ubuntu), 72–73
 POP configuration and, 325–326
 set-up, 325–326
timeout values (SELinux), CD169
TimeVault, CD200, CD201
title slides (Impress), 700–701
TLS (Transport Layer Security), 1060
ToDo items, managing, 416
toolbars
 Calc, 672
 Draw, 733
 Evolution, 328–329
 Gnumeric, 681–682
 for HTML e-mail, 405
 Impress, 699
 Inkscape, 737–738
 Karbon14, 740–741
 Krita, 724–725
 KSpread, 689
 KWord, 636
tools. *See also* command-line tools; graphical tools
 Adept tools. *See* Adept Package Manager
 basic user-level tools, 931
 gconftool-2, 181–182
 GNOME configuration tools, 180–185
 gnome-nettool package, 1080
 GTK, 680
 for package support, 67
 parted tool, 1033
 qtparted tool, 1033
 SELinux, CD169–170
 Software Preferences tool, 940
 Software Sources tool, 936–942
 tracker tool (GNOME), 265
 wireless tools packages, 1102
Torcs arcade game, 857
torrent, defined, CD58
torrents, finding, CD62–63
Totem application, playing DVDs with, 783–784
Tracker desktop search application (GNOME), 162–166
transferring files. *See* file sharing
translation and localization, 8
translucency, defined, 176
Transmission on Ubuntu, CD59–60
transparency, defined, 176
Transparent 3D theme (AWN), 195–196
Transport Layer Security (TLS), 1060
transport protocol, selecting (Gmail), 343–345
Trash applet (KDE 3), 208
Tresys Bedrock, CD169
Trolltech, 204
troubleshooting
 Apache, CD285–286
 boot problems, 1030–1032
 devices, 1045–1046
 DHCP servers, CD319–320
 DNS servers, CD335–337
 installation problems, 83–84
 profile migration, 478–479
 remote printing, CD350–352
 when adding drives, 1028
tsclient for terminal server connections, CD22–25
TTL (Time To Live), zone files and, CD327–329
tutorials
 digital graphics, 709
 Firefox, 429
 GIMP, 721

k3b, 721
LaTeX, 615
Linux Shell Scripting Tutorial, 566
OpenOffice.org, 633
PenDriveLinux.com, 57
Scribus, 647–656
spreadsheets, 660–669
Writer, 633
Tux Racer arcade game, 856
Twitter Microblogging (KDE 4), 296

U

ubufox package, 423
Ubuntu. *See also* startup process (Ubuntu/Kubuntu)
Add/Remove menu item, 927–929
advantages of, 455
basic tools, 931
benefits of, 11–12
booting, 45–47
CDs available for, 18–19
commercial support for, 23–24
community, 925
configuring sound on, 750–760
Debian Project and, 10–11
digital cameras on, 873–878
documentation, 22–23
games on, 825–830
information references, 24–25
installing. *See* configuring at installation (Ubuntu); installing on special-purpose systems (Ubuntu/Kubuntu); installing Ubuntu/Kubuntu; Ubuntu Server Install CD options
installing KDE on, 271
integrating Mac OS X printing with, CD347–350
integrating Windows printing with, CD346–347
migration, and installing, 489–491
migration, checking system before, 477–478
Network Manager on. *See* Network Manager on Ubuntu
overview, 6–7
printers, adding, 995, 997–1002
printers, configuring, 995–997
readahead and, 920
recognition of digital devices on, 860–865
release schedule, 8–9
repositories and components. *See* repositories and components
scanners and. *See* scanners
services, configuring at startup, 919
shutdown process, 921–922
spreadsheets, 659–660
startup. *See* startup process (Ubuntu/Kubuntu)
support for, 19–24
Ubuntu Manifesto, 7–8
update and maintenance commitments, 9–10
version, selecting, 28
versions for end users, 12–15
versions for server users, 15–16
versions for specialty platforms, 16
Ubuntu Desktop CD, 65
Ubuntu Mobile Internet Device (MID Edition), 16
Ubuntu Server Install CD, 18–19
Ubuntu Server Install CD options, 67–93
basics, 67–68
booting, 93
partition layout, specifying, 84–93
problems, resolving, 83–84
server, installing. *See* configuring at installation (Ubuntu)
Ubuntu Update Manager, 970–972
ubuntu-restricted-extras package, 794
ubuntu-vm-builder script, CD254
UCE (Unsolicited Commercial E-mail) defined, CD300
udhcpd, CD310
ufw firewall package, 1085
UltraSPARC systems, 28
umasks (user masks), 558
uml-utilities package, CD228
undeleting, e-mail messages, 334
undo (vi), 578
Uniform Resource Identifiers (URIs), CD278
Universal Resource Names (URNs), CD279

Unix
editing text on, 569–571
Linux/UNIX backup utilities, CD179
printing history, CD339–340
shells and, 542–543
Unsolicited Commercial E-mail (UCE), defined, CD300
Update Manager
basics, 931
using, 970–972
updates. *See also* Adept Updater, on KDE 4
searches for with Adept Manager, 977–979, 986–987
Software Sources tool and, 939–940
specifying update mechanisms, 80
updating Thunderbird, 462
upgrading systems
with apt-get, 951–952
with Synaptic, 962
uplink ports, 1052
Upstart system initialization, 913–916
URIs (Uniform Resource Identifiers), vs. URLs, CD279
URLs
svn in Firefox, CD106
vs. URIs, CD279
URNs (Universal Resource Names), CD279
USB
Flash drives, reformatting, 59
printers, adding, 995, 1012
USB sticks, creating bootable, 57–60
user interfaces
aptitude and, 955–956
Firefox, 424–428
user masks (umasks), 558
user networking, with KVM, CD258–259
usermode command, CD110
users
access to printers, 1008–1009
authorized for rsync backups, CD192
importing mail and, 462–463
KDE 3, managing, CD116–119
KDE 4, CD121–123
NFS authentication, CD360
passwords, creating at installation, 77–78
permissions and, 118–119
Ubuntu, managing, CD110–114
Ubuntu, versions for end users, 12–15
Ubuntu versions for server users, 15–16
user accounts, creating, 75, 76
usernames/passwords (Firefox), 433–434
utilities. *See also* tools
AppArmor, CD161–162
apt-cache, 946, 947
apt-file, 948, 949
aptitude, vs. dselect, 954
convert, 718
cpio, CD179
dd backup, CD179
deborphan, 991–992
dlocate, 949
dselect, 931, 932, 954
dump/restore, CD179–180
fdisk, -l (list) option, 111
fsck, 109
GNOME DHCP configuration utility, CD315
graphical Disks utility (Ubuntu), 1032
gufw graphical configuration utility, 1085–1090
Linux/UNIX backup utilities, CD179
readoe, 473
testparm, CD378–379
uml-utilities package, CD228

V

values
calculated (KSpread), 690
currency in spreadsheets, 675, 676, 683, 684–686, 693
pasting into cells (spreadsheets), 664–665
timeout values (SELinux), CD169
variables
definitions (AppArmor), CD158
emacs internal variables, 593–594
environment variables, 566–567
vCards, 356
.vcf file extensions, 356

vector graphics
 defined, 709
 tools for Linux, 735–736
Venema, Witse, CD289, CD305
verifying, backups, CD176–177
Veritas's Backup Exec software, CD178
vi, 571–583
 basics, 571
 to compose e-mail messages, 408
 cutting, copying and pasting in, 576–577
 deleting and changing text in, 575–576
 vs. emacs, 571
 graphical versions of, 582
 inserting text in, 574
 moving around in, 574–575
 multiple windows in, 579–580
 resources, 583
 starting and exiting, 572
 text search and replace in, 577–578
 undoing changes in, 578
 vs. vim, 572
 vim, customizing, 580–581
video, digital. *See* digital audio/video (Kubuntu); digital audio/video (Ubuntu)
video burn, 147
video games (Ubuntu), 829–830
viewport, defined (Compiz), 178
vim
 customizing, 580–581
 resources, 583
 vs. vi, 572
 vim-tiny, 572
.vimrc file, 581
vinagre client, CD7
vino VNC server, configuring (Ubuntu), CD7–9
virtual bridges, CD259–260
virtual desktops, defined (Compiz), 178
virtual folders, defined (Evolution), 336
virtual machines
 boot/installation problems, fixing, CD233–234
 creating, CD244–248
 emulation software for Linux, CD208–209
 hardware support for, CD207
 hardware virtual machine DomainU, creating, CD271–272
 installing operating systems in, CD220–223, CD255–256
 KVM. *See* KVM (Kernel-based Virtual Machines)
 managing, CD209
 overview, CD203–204
 QEMU CPU emulator, CD273–275
 reasons for, CD204–205
 resources, CD210
 selecting virtualization technology, CD208
 starting and stopping, CD224–226
 types of, CD205–206
 Virtual Machine Manager, CD209
 VirtualBox. *See* VirtualBox
 VMware server. *See* VMware server
 VMware virtual machine, starting, CD249
 Xen technology. *See* Xen technology
Virtual Network Computing. *See* VNC (Virtual Network Computing)
Virtual private networks (VPNs), configuring connections (Ubuntu), 1062–1067
VirtualBox
 creating virtual machines in, CD213–220
 installing, CD212–213
 installing Guest Additions, CD224
 networking and, CD227–232
 overview, CD211–212
 VM access, CD232
viruses
 detecting (Kontact), 411–414
 scanning (Postfix), CD300–305
VLC Media Player, 774
VM access, remapping ports for (VirtualBox), CD232
VMware server
 downloading and installing, CD235–242
 kernel modules, building, CD240
 overview, CD234–235
 starting, CD242–243
VNC (Virtual Network Computing)
 connection to other systems with, CD7

servers, configuring, CD7–12
servers/clients for Mac OS X, CD12
servers/clients for Windows, CD13
Volume Control applet (GNOME), 145
volume groups, defined, 108
VPN connections, configuring (Ubuntu), 1062–1067

W

Water Effect plug-in, 192
wear-leveling (Flash drives), 57
weather
displaying information, 367
Weather Report applet (GNOME), 145
Web Distributed Authoring and Versioning (WebDAV), CD100
Web servers, CD280–286
Apache, basics, CD280
Apache, configuring, CD282–284
Apache, troubleshooting, CD285–286
Apache 2 file locations, CD280–281
basics, CD279–280
installing, CD280
resources, CD286
URL vs. URI, CD279
World Wide Web, CD277–279
Web shortcuts, 260–261
Web sites, bookmarks, storing on, 482
Web sites for downloading
Adium X, 496
Advanced Shell Scripting Guide, 566
Atomic Tanks, 852
bcm43xx-fwcutter, 1107
Beryl Emerald themes, 198
Bogofilter, 341
Bonjour software, 499
Chicken of the VNC (client), CD12
Dragon player, 786
FileZilla FTP client, CD44
Firefox add-ons and themes, 442
FreePOPs!, 347
Frozen Bubble arcade game, 852
Gadu-Gadu, 499
games for kids, 839
GNOME themes, 155
gnubg, 835
Groupwise software, 499
gtk-gnutella, CD58
installation CDs, 66
IRC clients, 526
Jduplicate, 851
JuK media player, 784
Kaffeine media player, 784
KBear FTP client, CD43
Kerberos, CD142
KFTPgrabber FTP client, CD43
Kubuntu examples folder, 47
MailMan, 357
MAME, CD208
MID Edition, 16
MySpace Messenger, 500
Noatun media player, 784
Parallel Realities (Linux games), 855
Pidgin IM client, 496
Plasma Themes, 311
plug-ins (Firefox), 453
pop3check application, CD93
PPD files, 1000
printer drivers, 999
QQ software, 500
Qt software, 204
Quackle, 851
rdesktop package, CD20
Samba software, CD45
Sametime software, 500
Scribus, 646
Separate+ plug-in, 714
Simple Instant Messenger, 500
SpamAssassin, 341
Thunderbird, 348, 460
TightVNC, CD12
Totem, 783
Tracker, 162, 265
Transparent 3D, 195
tsclient application, CD22

Ubuntu Desktop CDs, 29
VMware-Any-Any package, CD240
WebCal calendars, 367
WINE software, 884
Wubi, 60
Yahoo! Mail Plus package, 349
Zephyr protocol (IM), 500
Web sites for further information, 16
anti-virus software, 411
Apache, CD286
AppArmor, CD167
aRts, 773
assistive technologies and GNOME, 156
Avant Window Navigator, 192
Avant Window Navigator Wiki, 197
AWN applets, 194
backup software, CD178, CD198, CD200
backuppc, CD188
BackupPC package, CD188
BIND, CD337
BitTorrent, CD63
Bogofilter, 411
bookmarks, 482
Bootchart, 916
Compiz Fusion project, 180, 185, 192
Cool-Streams, 818
CUPS, CD352
custom fields (Kontact), 389
CVS repositories, CD100
Debian Project, 10
Decibel (KDE 4), 276
Disk Manager, 1032
DistroWatch, 24
DNS servers, CD337
Dolphin file manager, 258
Draw, 735
Dropbear, CD162
Eclipse, CD99
Edubuntu, 13
emacs, 593, 595
e-mail, free accounts, 343
e-mail access, CD299
Evolution, 370
Fedora, CD169
FFMPEG project, 803
Firefox, 429
Fluendo, 776
Fluxbuntu, 13
Free Internet Chess Server, 843
FreeRADIUS server, 1060
The Fridge, 22
Geubuntu, 13
GIMP, 721
Gmail, 346, 347
GNOME Format, 1032
Gobuntu, 13
graphics, 709
GraphicsMagick package, 723
GRUB, 909
GStreamer, 773, 775
Heimdal client/server implementation, CD142
IDEs, CD81
IETF's RFC archive, CD308
IM accounts, 497
IM application packages, 531–532
IM clients, 496
IM inside emacs, 495
IMAP configuration, 347
init packages, 911
Inkscape, 738
installation CDs, 66
International Go federation, 847
Internet Scrabble Club, 851
IP addresses, 1052
IRC, 525, 531
JACK Audio Connection Kit, 773
junk e-mail, 341
k3b, 792
Karbon14, 741
KDE, 203, 250
KOffice, 633
Kopete IM client, 525
KPhotoAlbum, 879
krfb VNC server, CD9

Krita, 729–730
Kubuntu, 14, CD170
KVM, CD252–253
KWord, 645
LaTeX, 615
Launchpad, Canonical, 15
LILO, 907
Linux equivalents to Windows applications, 491
Linux for You magazine, CD108
Linux package formats, 990
Linux printer support, 994
Linux Standard Base project, 106
Linux Terminal Server, CD28
LTSP, 99
Lyx, 613
make program, CD70
MPlayer Movie Player, 774
MTAs, CD289, CD290
NAS, 773
ndiswrapper, 1104, 1108
Network Manager, 1053
NFS, CD370
NoMachine NX Server, CD27
Novell Core Protocol, 111
Open Clip Art Library, 706
OpenBSD project, CD5
OpenGL FAQ, 176
OpenOffice.org, 615
OpenSSH server, CD5–6
PenDriveLinux.com, 57
Phonon (KDE 4), 276
Phonon Multimedia API, 773
Pidgin, 509
PKI, CD5
Planet Ubuntu blog, 20
plasmoids, 291
POP configuration, 346
port forwarding, CD58–59
Postfix, CD297
Postfix mail servers, CD305
PulseAudio sound server, 773
PulseAudio wiki, 779
PVM clustering software, CD205
QEMU CPU emulator, CD273–275
RBAC, CD166–167
RealVNC, CD7
Red Hat, CD169
RFCs, CD322
Samba server, CD380
scanners, 1018
Scribus, 646, 656
security mechanisms, 1079
SELinux, CD169–170
shells on Unix, 543
Shoutcast, 818
SlashDot, 297
Solid (KDE 4), 276
SourceForge site, CD20
Sox project, 803
Strigi desktop search application, 264, 267
Subversion, CD100, CD107–108
TeX, 615
TeX Live, 606
TeXmacs package, 614
text editing, 571, 603
THE, 603
TightVNC, CD7
Ubuntu, CD170
Ubuntu CDs, 19
Ubuntu Christian Edition, 14
Ubuntu documentation, CD132
Ubuntu forums, 20–21
Ubuntu Linux, 23, 24–25
Ubuntu mailing lists, 22
Ubuntu Manifesto, 7
Ubuntu marketplace, 24
Ubuntu Muslim Edition, 14
Ubuntu repositories, 933, 934
Ubuntu Studio, 15
Upstart, 916
vim, 581
vino VNC server, CD7
Virtual Machine Manager, CD209
VirtualBox, CD211

virtualization, CD210
VLC Media Player, 774
VMware, CD234–235
WEP keys, packages for cracking, 1093
wvdial, 1079
X Window system, 124–125
Xen technology, CD272–273
XenoServers project, CD261
XenSource, Inc., CD261
Xine Multimedia player, 774
X-KDE-Solid-Predicate test, 872
Xubuntu, 15
Web-based e-mail. *See* Gmail, configuring; Gmail, setup for reading
Web/Browser settings, migrating to Linux
basics, 475–476
bookmarks. *See* bookmarks, migrating to Linux
from Firefox to Firefox, 476–479
from Internet Explorer/Safari to Firefox, 479–481
WebDAV (Web Distributed Authoring and Versioning), CD100
WEP (Wired Equivalent Privacy)
basics, 1093
WEP 128-bit Passphrase, 1060
WEP 40/128-bit Key, 1060
whitelisting, defined, CD301
whitespace, around frames (KWord), 643
Wicd, 1094–1102
configuring, 1100–1102
installing, 1095
using, 1096–1099
Wicd package, 1061
wide-dhcpv6-server, CD310
widgets, desktop (KDE 4)
adding, 291–293, 296–297
default widgets, 293–296
overview, 290–291
Wi-Fi Protected Access (WPA). *See* WPA (Wi-Fi Protected Access)
wildcards
bash shell and, 562–563
in dpkg specifications, 944
window managers
changing (KDE 3), 268–270
GNOME desktop and, 170–174
Window Management desktop effects (KDE 4), 307
windows
emacs (vi and vim), 588
Firefox, 424–425
multiple in emacs, 592–593
multiple in vi, 579–580
window captures (GIMP), 714
window decorations (KDE 4), 278–280
Windows
file sharing, CD371–372
networking drivers, 1104–1108
Pidgin IM client, 496
printing, integrating with Ubuntu, CD346–347
users for backups, CD194
VNC servers/clients for, CD13
Wubi, installing Ubuntu/Kubuntu on, 60–63
Windows Live Messenger, 497
Windows Printer via SAMBA printer connection, 998, 1013–1014
Windows shares
accessing, CD45–46
connecting from command line, CD53
connecting from GNOME desktop, CD48–50
connecting from Kubuntu, CD50–53
connecting in Dolphin, CD52–53
connecting in Konqueror, CD50–52
locating, CD46–47
mounting at boot time, CD47–48
Windows systems, migrating from, 455–491
basics, 455–457
e-mail, migrating to Linux. *See* e-mail, migrating to Linux; Thunderbird, migrating e-mail via
personal data, backing up, 457–458
Web/Browser settings. *See* Web/Browser settings, migrating to Linux
Windows application data, 489–491
Windows Terminal Services, CD20–22
WINE software, 884
Win-Modems, 1079

wired connections, configuring (Ubuntu), 1056–1057
Wired Equivalent Privacy. *See* WEP (Wired Equivalent Privacy)
wired/wireless connections
 configuring on KDE 3, 1071–1074
 configuring on KDE 4, 1074–1076
wireless connections, configuring (Ubuntu), 1057–1061
wireless networking, 1091–1108
 command-line tools for, 1102–1103
 defined, 1092
 Wicd. *See* Wicd
 Windows networking drivers, 1104–1108
 wireless technologies overview, 1091–1094
wireless security on KDE 3, 1072
wizards
 Anti-Spam Wizard, 411, 412
 Anti-Virus Wizard, 411
 Document Converter Wizard, 617, 631, 670, 678–680, 702
 document wizard (Karbon14), 739–740
 documents, creating with, 622–625
 Image Wizard (Krita), 723–724, 727
 Import Wizard (Evolution), 466, 467–469
 Import Wizard (Firefox), 486
 Import/Export Wizard (Internet Explorer), 482–484
 for importing spreadsheets, 678–680
 Impress Presentation Wizard, 696–698
 Letter Wizard, 622–624
 Printer Wizard, 1002–1010
 Safari Bookmark Exporter, 484–485
 for starting KSpread, 688–689
Wobbly Windows plug-in (Compiz), 179, 181–185, 192
WordBiz (game), 851–852
work hours/day/week (Calendar), 359–360
workgroups (Samba), CD374–375
workspace, defined (Compiz), 178
World Wide Web, CD277–279
Wormux arcade game, 857
WPA (Wi-Fi Protected Access)
 basics, 1094
 and WPA2 Enterprise/Personal, 1060–1061
Writer. *See* OpenOffice.org Writer
Wubi (Windows-based Ubuntu Installer), 60–63
wvdial, 1079

X

X MultiMedia System. *See* XMMS (X MultiMedia System)
X servers/X Window servers, 177
X Window system, 123–124, 548–551
x86
 installation, 74
 system boots, 66
Xara Xtreme, 743
xboard chess engine, 844
XChat UNI X/Linux IRC client, 528
XDR (eXternal Data Representation), CD357
XEmacs, 585–586, 945
Xen technology
 DomainU configuration files, creating, CD263–267
 information references, CD272–273
 installing, CD263
 overview, CD261–263
 paravirtualized domainU system, CD267–272
XenoServers project, CD261
Xgl server, 177
Xine Multimedia player, 774
xinetd daemon, CD105
XLS files, 491
XMMS (X MultiMedia System)
 audio player, 822
 building, CD70–71
 building SHN plug-in for, CD78–79
XSane Image Scanner software, 1018–1019
xterm
 basics, 548–551
 window (vi), 573
Xubuntu, 7, 15

Y

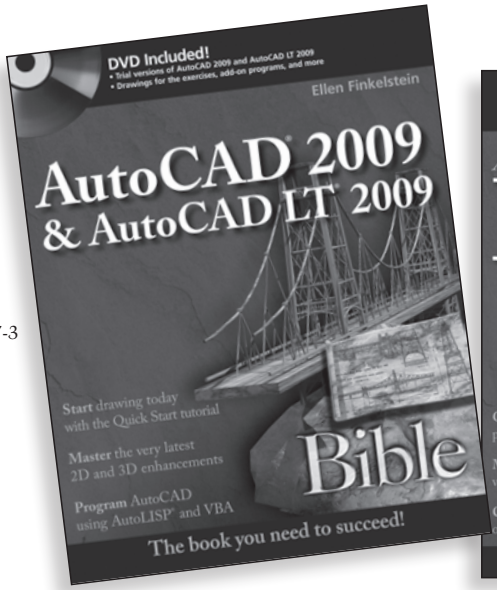
y (yank) command (vi), 576
Yahoo!, POP access to, 349
Yahoo Messenger software, 498
Yamakawa, Yoshinori, 714

Z

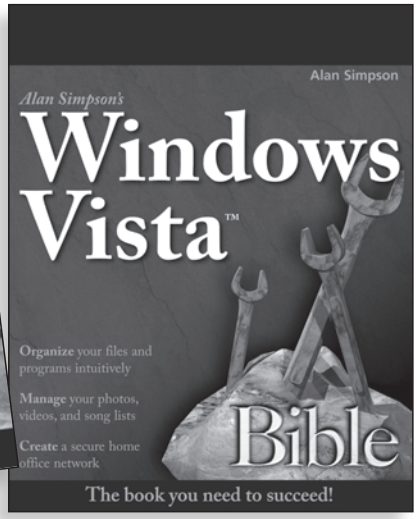
Zephyr protocol (IM), 500
.zip-format archive, for restoring files, CD196–197
zone files
 creating, CD327–331
 errors, CD335–337
 incorporating with BIND, CD332–333

The books you read to succeed.

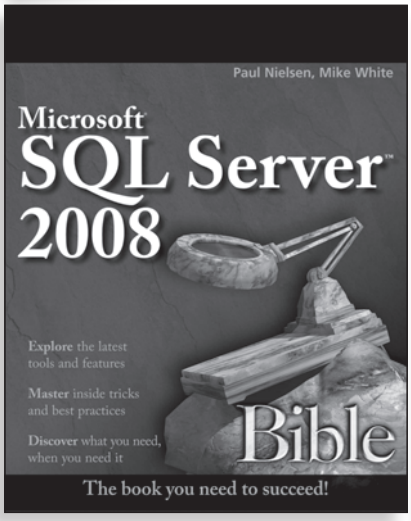
Get the most out of the latest software and leading-edge technologies with a Wiley Bible—your one-stop reference.



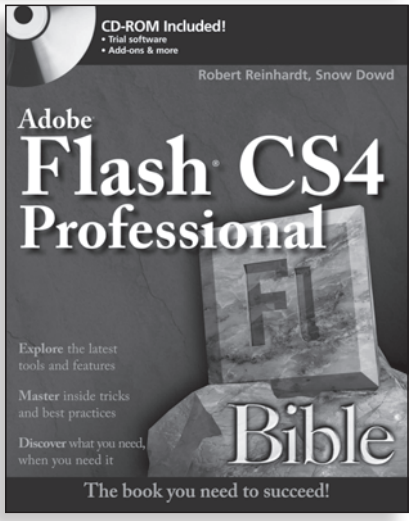
978-0-470-26017-3



978-0-470-04030-0



978-0-470-25704-3



978-0-470-37918-9

Available wherever books are sold.

Wiley and the Wiley logo are registered trademarks of John Wiley & Sons, Inc.
All other trademarks are the property of their respective owners.

WILEY
Now you know.
wiley.com

Part IV

Bonus Chapters

IN THIS PART

Bonus Chapter 1
Connecting to Other Systems

Bonus Chapter 2
File Transfer and Sharing

Bonus Chapter 3
Software Development on Ubuntu and Kubuntu

Bonus Chapter 4
Managing Users, Groups, Authentication, and Advanced Permissions

Bonus Chapter 5
Managing System Policies and Security

Bonus Chapter 6
Backing Up and Restoring Files

Bonus Chapter 7
Using Virtual Machines and Emulators

Bonus Chapter 8
Setting Up a Web Server

Bonus Chapter 9
Setting Up a Mail Server

Bonus Chapter 10
Setting Up a DHCP Server

Bonus Chapter 11
Setting Up a DNS Server

Bonus Chapter 12
Setting Up a Print Server

Bonus Chapter 13
Setting Up an NFS Server

Bonus Chapter 14
Setting Up a Samba Server

Bonus Chapter 1

Connecting to Other Systems

Not too long ago, networking meant either talking with groups of friends or sending e-mail and surfing the Web from your home computer. Today's multi-computer households can provide much more than that, often featuring home or SOHO (small office/home office) networks of computers that share files and peripherals (such as printers and scanners); they are often centrally administered, and provide access to multiple people who want or need it. Although networked access to shared resources such as file servers and printers has been common in enterprise and home Windows networking environments for a long time, this is now a common requirement for the modern, enlightened household.

This chapter discusses various command-line and graphical solutions provided by Ubuntu Linux for directly connecting to remote systems, either by logging on to them directly or by obtaining remote, graphical access to a desktop session. This chapter doesn't explore connecting to remote file servers and file sharing — that's discussed in Bonus Chapter 2, "File and Transfer and Sharing on Ubuntu."

IN THIS CHAPTER

Establishing secure connections to other systems

Enabling incoming connections to your machine

Accessing remote systems graphically

Establishing Secure Connections to Other Systems

The traditional UNIX application that establishes a terminal connection (a command-line, login connection) to a remote machine over the network is known as *telnet*. The `telnet` application connects to a `telnet` daemon running on a remote system and creates a login session once the proper authentication credentials have been exchanged. Unfortunately, authentication information and all other data exchanged between a `telnet` client and server is exchanged in the clear, meaning that it is readily visible to anyone with access to your network, a packet sniffer, and sufficient curiosity or other motivation. There are

ways around this, of course: You can use `telnet` through an encrypted virtual private network (VPN) or tunnel, or install a version of `telnet` that uses Kerberos authentication, but why cling to an old application and protocol when there are newer, more secure alternatives available? (If you're still fascinated with `telnet`, search the Web for information about Linux VPNs or see Bonus Chapter 4 for information about installing Kerberos on your system, which includes a Kerberized version of `telnet`.)

TIP **The `telnet` client is installed as a part of a default Ubuntu installation because you may need the `telnet` client to connect to older systems or systems that simply do not support more sophisticated remote connections, but a `telnet` server is not officially supported by Ubuntu. If you must, you can still explicitly install the `telnet` server — known as `telnetd` — from a "universe" repository by using your favorite package management software.**

If you're an old-time UNIX fan, you may remember the Berkeley r-command suite, which provided commands such as `rsh` and `rcp`. These commands had the same security problems as the older `telnet` commands, namely that they transmitted passwords and other information in the clear. Although these commands appear to be available on Ubuntu Linux systems, they are mapped to the corresponding SSH commands (`ssh` and `scp`, respectively) using the *alternatives* system that is discussed in Chapter 23, "Adding, Removing, and Updating Software."

Logging in to Remote Systems Using ssh

The newer, more modern and secure replacement for `telnet` is `ssh`, the secure shell application, which uses SSH, the Secure Shell protocol. The `ssh` application is installed as part of a default Ubuntu Linux installation, along with other applications that use the same protocol and authentication information, and shares much of the same command-line syntax, such as `sftp` (secure FTP) and `scp` (secure copy). Using `ssh`, you can establish a remote login session on any machine that is running an `ssh` daemon. An `ssh` daemon must be running on the remote system. If you try to connect to the remote system and see a message like the following, either an `ssh` daemon is not running on that system or it is running on some port other than the default port 22:

```
$ ssh ulaptop
ssh: connect to host ulaptop port 22: Connection refused
```

The remainder of this section discusses how to use `ssh` to connect to remote systems that are actually running an SSH daemon. For information on how to enable the SSH daemon on Ubuntu and Kubuntu systems, see the next section.

Using `ssh` to connect to a remote system as the current user is as simple as typing the `ssh` *`hostname`* command, where *`hostname`* is the name or IP address of the host to which you want to connect. The first time that you attempt to connect to a remote system that you have never connected to before, you will see a message like the following:

```
$ ssh ulaptop
The authenticity of host 'ulaptop (192.168.6.90)' can't be established.
RSA key fingerprint is 07:e6:3a:50:4b:6d:e6:d8:f1:80:c6:b2:da:02:a3:da.
Are you sure you want to continue connecting (yes/no)?
```

To continue connecting to this system, type **`yes`** and press Return. The `ssh` utility adds the RSA key for the remote system to its list of known hosts (to prevent another machine from impersonating the remote machine in the future) and prompts you for your password, as in the following example:

```
Warning: Permanently added 'ulaptop' (RSA) to the list of known hosts.
wvh@ulaptop's password:
```

NOTE **If you're curious, RSA stands for Rivest, Shamir, and Adleman, the developers of the type of public key encryption technology that is used by default by the OpenSSH implementation of the SSH protocol. For more information about public key infrastructure (PKI) and related encryption technologies, see a Web site such as Wikipedia's PKI information at** `http://en.wikipedia.org/wiki/Public_key_infrastructure`.

At the password prompt, enter your password on the remote system and press Return. Assuming that you typed the correct password, you'll see a successful login message from the remote system, and that system's prompt will display, as in the following example:

```
Linux ubuntu32 2.6.24-19-386 #1 SMP Fri Jul 11 16:43:33 UTC 2008

The programs included with the Ubuntu system are free software;
the exact distribution terms for each program are described in the
individual files in /usr/share/doc/*/copyright.

Ubuntu comes with ABSOLUTELY NO WARRANTY, to the extent permitted
   by applicable law.
Last login: Wed May 1 20:15:16 2008
wvh@ubuntu32:~$
```

Once you've established an `ssh` connection to a remote machine, you can do anything from this login prompt that you could do from a physical connection to that machine, including starting graphical applications that display on your local system if you set the `DISPLAY` environment variable correctly, as discussed in the section of Chapter 5 entitled "What's a Desktop? Graphical Environments for Linux." When you're done working on the remote system, you can log out by pressing Ctrl+D or by typing the `exit` command.

All of this is well and good assuming that you have the same login name on all of the systems that you want to connect to, or that you always want to connect as yourself. Luckily, the `ssh` command provides an easy way to connect to another system as another user by using a command such as `ssh` *`user@host`*, where *`user`* is the name of the user that you want to log in as, and *`host`* is the name of the host to which you want to connect as that user. For example, the command `ssh joeuser@ulaptop` would use `ssh` to connect to the remote host `ulaptop` as the user `joeuser`. Assuming that you know this user's password, you can log in there as the specified user and perform any actions that the user is authorized to do.

The `ssh` command has many other options, few of which you'll probably ever need to use. For complete information on the `ssh` command, see the online reference information for the command that is available through the `man ssh` command.

Enabling Incoming SSH Connections to Your System

As mentioned previously in this chapter, an SSH daemon must be running on your system if you want to be able to connect to it remotely using any command that requires the SSH protocol. The standard version of the SSH server that is available for Ubuntu and Kubuntu systems is the OpenSSH server (`www.openssh.org`), an open source SSH server that is actively maintained by the people who also support the OpenBSD project (`www.openbsd.org`), an open source UNIX-like operating system (inspired by the original Berkeley UNIX distributions) that is well known for its focus on system security. The OpenSSH server is not installed by default as part of a standard Ubuntu or Kubuntu installation but (like everything else in Ubuntu/Kubuntu) is easy enough to add using any of the package management tools described in Chapter 23.

To install the OpenSSH server, use your favorite software installation package to install the `openssh-server` package. I prefer using the graphical Synaptic Package Manager on Ubuntu systems and the Adept Package Manager on Kubuntu systems, but feel free to use the command-line package management tools if you prefer.

NOTE **See Chapter 23, "Adding, Removing, and Updating Software," for more detailed information on installing applications on your Ubuntu or Kubuntu system.**

As part of installing the OpenSSH server, Ubuntu installs a control for it in the administrative Services application, which you can access by selecting Services from the System ➪ Administration menu. As shown in Figure B1-1, the Remote shell server is automatically enabled when you install it.

FIGURE B1-1

Verifying the state of the SSH server in the Services application



Should you ever want to stop enabling remote users from attempting to log into an Ubuntu system, click Unlock to display an authentication dialog in which you must specify your password. You can then deselect the checkbox beside the Remote shell server entry in this application and click Close to terminate the services configuration application and deactivate the OpenSSH server.

Kubuntu systems provide a similar administrative application for enabling and disabling an OpenSSH server (as well as other servers and the services that they provide). On KDE 3-based Kubuntu systems, select K Menu ➪ System Settings, and click the Advanced tab. Double-click the System Services icon, and scroll down in the list displayed in the left pane to see the current status of the ssh service that is associated with the OpenSSH server. To disable the SSH server, click Administrator mode and supply your password in the dialog that displays. You can then select the ssh entry in the left pane and click Stop to temporarily stop the OpenSSH server. You can also de-select the Start during boot checkbox if you do not want the OpenSSH server to start automatically the next time that you boot your system. Clicking Close to close the System Services application will save your changes.

TIP **The `dropbear` package is an attractive alternative to the OpenSSH server, supporting most of the features of the OpenSSH server but using substantially fewer system resources. I typically use `dropbear` on laptops, netbooks (small-format, network-oriented portable computers), mobile Internet devices (tablet-like portable computers designed for Internet access), and any other limited-resource systems where I want to support secure, incoming SSH connections. `Dropbear` is started and stopped by the `/etc/init.d/dropbear` startup script, and is not integrated with the Services window shown in Figure B1-1.**

Connecting to Other Systems Using VNC

If you have multiple computer systems on a single network, getting up and going to the console of a specific machine to run a program there quickly becomes old. You'd think there would be a way that a graphical computer system could connect to another graphical computer system over the network and display what you would see on that machine's screen inside a window on the first machine. Actually, there is, and it's included in the standard Ubuntu and Kubuntu distribution by default.

Virtual Network Computing (VNC) makes it easy for you to work on multiple graphical computer systems by enabling you to export their graphical displays across a network to another machine. VNC is a cross-platform technology originally developed by Olivetti Research Labs in Cambridge, England, which was later acquired by AT&T. A VNC server runs on one computer system and exports a system desktop that can be accessed by a VNC client running on another system. VNC servers are typically password-protected, which makes VNC an optimal environment for accessing and running graphical administrative and monitoring applications remotely. VNC servers export their virtual displays via network ports starting at 5900 plus the number of the graphical displays that they are exporting. For example, a VNC server running on the X Window system display `:1` will use port 5901.

VNC is released under the GPL, and many of the original VNC developers now work for a company called RealVNC (`www.realvnc.com`), which distributes and supports a commercial VNC implementation but also makes the source code for its version of the VNC server and client freely available, as required by the GPL. The latest generation of their VNC client is named `vnc4viewer`, and is available in the Ubuntu repositories but is not installed by default. Another extremely popular VNC implementation is TightVNC (`www.tightvnc.com`), an optimized VNC client and server, which is also available for Ubuntu and Kubuntu systems but is also not installed by default.

Both Ubuntu and Kubuntu systems provide graphical VNC clients that are installed by default as part of a standard desktop installation. The Ubuntu client that is installed by default is `vinagre`, and the Kubuntu client that is installed by default is `krdc`. Don't worry — menu entries are provided for both of these so that you don't have to try to remember their names. Using these applications is discussed later in this chapter.

Configuring a VNC Server

Before trying to use a VNC client to connect to a remote system, you must have a VNC server to connect to. This section explains how to enable the VNC servers that are installed by default on Ubuntu and Kubuntu systems.

TIP **Regardless of the platform on which you are running a VNC server, you must make sure that the network ports that it uses are not being blocked by other network software. For example, if your system does packet filtering or your network uses a firewall, you must make sure that you do not block ports 590*x* (used to export VNC server displays), 6000 (used to communicate with Linux or Unix X Window system servers), or 580*x* (used to communicate with a VNC server over the Web).**

Enabling and Configuring Ubuntu's vino VNC Server

Connecting to an Ubuntu system on your local network is quite simple — the software that you need to connect to the current graphical console is already installed. This is a GNOME-specific VNC server called vino (the actual executable is called `vino-server`), which was created by Mark McLoughlin. Little other information about vino is available beyond the initial check-in document at `http://cvs.gnome.org/bonsai/cvsblame.cgi?file=vino/docs/remote-desktop.txt` and the minimal information in `/usr/share/doc/vino` on your Ubuntu system. However (luckily), in my experience it "just works," and therefore I guess we don't really need to know much more.

NOTE **Vino works differently from traditional VNC servers, in that once it is enabled, the GNOME session manager automatically starts it for you whenever you log in. It can therefore export only the display of a user that is currently logged in, unlike more standard VNC servers that typically display a system's login screen and therefore support remote access to your system by random users. I discuss the high points of more traditional VNC servers later in this section, in the sidebar "How Traditional VNC Servers Work."**

To make sure that the vino VNC server is enabled for you, select the System ➪ Preferences ➪ Remote Desktop menu item, which displays the dialog shown in Figure B1-2. To enable remote access to your desktop, select the Sharing section's Allow other users to view your desktop checkbox. This enables remote users to view the desktop but not to actually do anything there — you must select the Allow other users to control your desktop checkbox to enable a remote VNC client to control the mouse and use keyboard shortcuts.

When configuring the vino server, always set a remote desktop password in the Security section of the dialog shown in Figure B1-2. To do this, select the Require the user to enter this password checkbox and enter a password. You can optionally select the Ask you for confirmation checkbox, which displays the dialog shown in Figure B1-3 on your desktop when someone tries to connect to it using a VNC client.

FIGURE B1-2

Configuring the vino VNC server on an Ubuntu system



FIGURE B1-3

A remote connection request on an Ubuntu desktop



In general, I don't recommend enabling this option, first, because it largely defeats the purpose of supporting remote VNC connections in the first place, and second, because access to your VNC server should always be password-protected anyway, so only people to whom you've given the password should be able to access your desktop via VNC.

Once you've enabled remote access using the Remote Desktop Preferences dialog, the GNOME desktop starts the `vino-server` process, and will start it automatically for you whenever you start a GNOME login session on this Ubuntu system.

Enabling and Configuring Kubuntu's krfb VNC Server

Connecting to a Kubuntu system on your local network is quite simple — the software that you need to connect to the current graphical console is already installed. This is a KDE-specific VNC server called krfb, which was created by Brad Hards. Like most KDE applications, krfb is actually documented, and you can find the documentation online at `http://docs.kde.org/stable/en/kdenetwork/krfb/index.html`.

NOTE **Like vino on GNOME-based Ubuntu systems, krfb works differently than traditional VNC servers, in that once it is enabled, the KDE session manager automatically starts it for you whenever you log in. It can therefore export only the display of a user that is currently logged in, unlike more standard VNC servers that typically display a system's login screen and therefore support remote access to your system by random users. I discuss the high points of more traditional VNC servers later in this section, in the sidebar "How Traditional VNC Servers Work."**

To make sure that the krfb VNC server is enabled for you, select the K Menu ➪ System Settings menu item to display the general KDE configuration application. Select the Sharing icon (which will display a warning if file-sharing is not enabled — more about that in Bonus Chapters 13 and 14). The sharing dialog displays. If necessary, select the Desktop Sharing icon in the left-hand panel to display the tabs shown in Figure B1-4.

As you can see, KDE's krfb configuration utility supports a transient remote access mechanism known as Invitations. This allows you to enable VNC access to a KDE-based system such as Kubuntu for a short period of time (one hour). You can either create a printable invitation that you can physically refer to (which can be handy for meetings) or send an invitation to a specified user via e-mail. These invitations contain access information and a unique, automatically generated password that will expire after an hour. Clicking the Create & Manage invitations button displays the Manage Invitation dialog that enables you to create both types of invitations and displays any invitations that are currently available. Figure B1-5 shows both the Manage Invitations dialog and an Email Invitation in the process of being created.

FIGURE B1-4

Configuring the krfb VNC server on a Kubuntu system

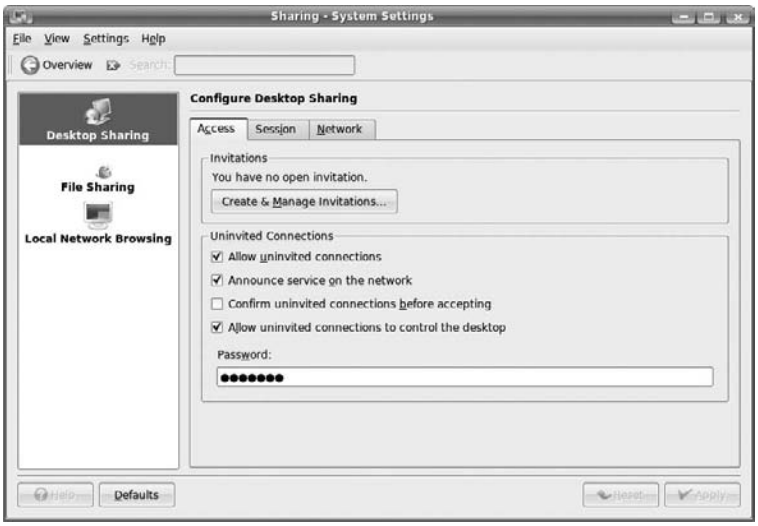
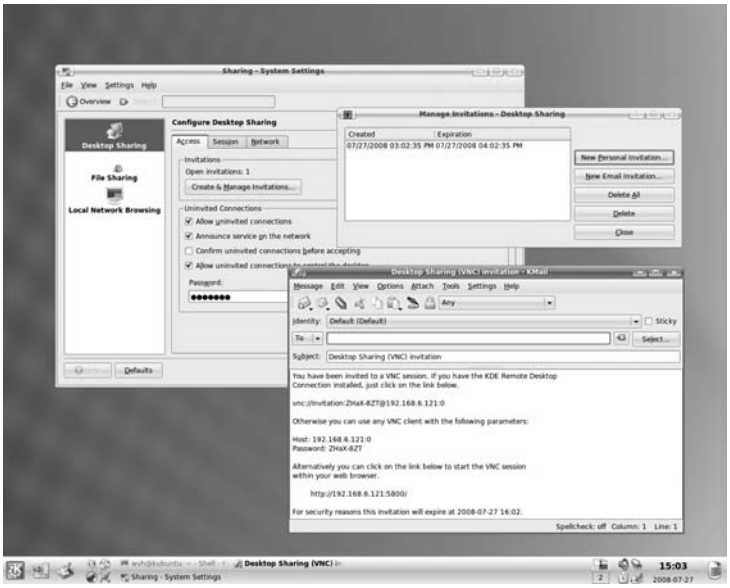


FIGURE B1-5

Sending a VNC invitation on a Kubuntu system



Invitations are great for customer support or teaching situations, where you want someone to be able to access your system's screen to help you or show you how to do something but you don't want to give out a global VNC password. Other VNC servers that I've run across don't do this, but they should.

The krfb server also supports standard access to your desktop using a generic password. To enable generic remote access to your desktop, select the Uninvited Connections section's Allow uninvited connections checkbox (see Figure B1-5 again). This enables remote users to view the desktop but not to actually do anything there — you must select the Allow uninvited connections to control the desktop checkbox (partially hidden in the figure) to enable a remote VNC client to control the mouse and use keyboard shortcuts. When configuring the krfb server, you should always set a remote desktop password at the bottom of the dialog shown in Figure B1-4.

You can optionally select the Confirm uninvited connections before accepting checkbox, which displays the Desktop Sharing dialog shown in Figure B1-6.

FIGURE B1-6

A remote connection request on a Kubuntu desktop



In general, I don't recommend enabling this option because it largely defeats the purpose of supporting remote VNC connections in the first place, and second, because access to your VNC server should always be password-protected. Krfb's support for short-term invitations makes it easy for you to grant temporary access to your system without having to share your VNS password.

Once you've enabled remote access using the Desktop Sharing dialog, the KDE desktop updates the `krfb` process, and the KDE session manager will start it automatically for you with your new settings whenever you start a new KDE login session on this Kubuntu system.

How Traditional VNC Servers Work

Typically, VNC servers are started by a system daemon such as inetd or xinetd, or by a shell or Perl script called `vncserver`, and export their own displays rather than the desktop of the currently logged-in user. These mechanisms provide some additional flexibility in your VNC server environment. VNC servers that are started by the inted or xinetd daemons in response to incoming connection requests can start multiple servers that provide different desktop sizes and resolutions, which is very convenient if your system has a 21-inch monitor and you want to connect to it from a laptop without having to scroll all over the place. If you start a VNC server using the `vncserver` script, you can also take advantage of user-specific startup scripts that identify the window manager and any X applications that your VNC server should start. The VNC server's startup script is usually the file `~/.vnc/xstartup`. If this directory and the startup file do not exist the first time that you start a VNC server, the directory is created and the startup script is cloned from the default X Window system startup file `/etc/X11/xinit/xinitrc`.

Exporting a graphical console that is independent of any physical one and using a startup script to do so enables traditional VNC servers to follow the same somewhat convoluted chain of X Window startup files that are normally used: `~/.xinitrc`, `~/.Xclients` on some systems, `/etc/X11/xinit/xinitrc`, `/etc/X11/Xsession`, and so on. These startup files can launch various desktop environments and window managers by using environment variable settings, and also provide a last resort of starting the `twm` window manager when no other window manager can be found.

VNC servers that follow this more traditional model and are available for Ubuntu and Kubuntu systems include the `vnc4server` and `tightvncserver` packages, which you can install using your favorite package management utility. However, if you install one of these on an Ubuntu system, you should un-install the `vino` VNC server. Similarly, if you install one of these on a Kubuntu system, you should un-install the `krfb` VNC server.

VNC Client and Server Software for Other Systems

If you still use other operating systems, you may want to use VNC to access them from your Ubuntu or Kubuntu desktop, or access your Ubuntu or Kubuntu desktop from the desktops of those other systems. This is no problem at all thanks to the fact that VNC is open source. The next two sections explain how to get VNC client and server software for popular platforms such as Mac OS X and Microsoft Windows.

VNC Servers and Clients for Mac OS X

A variety of VNC servers and clients are available for Mac OS X. Mac OS X has included a built-in VNC server since Mac OS X 10.4 (Tiger), but older Mac OS X releases required that you install one yourself. My favorite VNC server for pre-Tiger Mac OS X systems is the OSXvnc server, which is available at `http://sourceforge.net/projects/osxvnc`. You can download this as a Mac disk image (.dmg) file that will mount automatically and start the installer. The OSXvnc server displays its own configuration dialog when you start it — you configure the built-in VNC server using the System Preferences application's Sharing applet.

As far as VNC clients for Mac OS X go, my favorite is called "Chicken of the VNC." It's a small, fast VNC client with a funny name — and how can you beat that combination? The latest version of the Chicken of the VNC client is available from the project's SourceForge site at `http://sourceforge.net/projects/cotvnc`.

VNC Clients and Servers for Windows

TightVNC is a small, highly optimized VNC client and server. TightVNC makes efficient use of network bandwidth through some clever tricks using JPEG compression for the display and differentiating between local cursor movement and cursor movement that needs to be communicated back to the VNC server. TightVNC also provides security features that RealVNC does not, including automatic SSH tunneling for security purposes. Although RealVNC clients and servers are available for Windows, I really like TightVNC and therefore recommend using it. Both stable and bleeding edge (i.e., versions that are actively under development) versions of TightVNC are downloadable from `www.tightvnc.com/download.html` or from their SourceForge site at `http://sourceforge.net/projects/vnc-tight`.

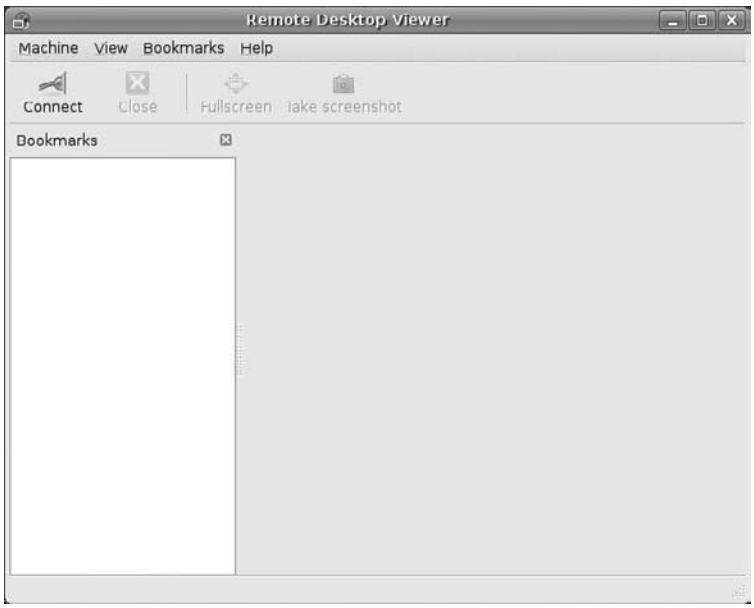
The downloadable EXE file from either of these sites provides both the VNC client and server. After downloading and installing the executables, you can start the VNC server from the standard menus. The VNC server's administrative dialog supports standard configuration options such as setting a password that you will need to use to connect to the VNC server, automatically starting the VNC server when you log in, and displaying an icon for the VNC server in the system tray when the VNC server is running. The icon is extremely convenient because it both lets you know when a server is running and also because you can get instant access to the configuration dialog by clicking on the icon.

Using Ubuntu's Remote Desktop Viewer

Once a VNC server is running on one of your systems, you can connect to it from any Ubuntu system by selecting the Applications ➪ Internet ➪ Remote Desktop Viewer menu command to start Ubuntu's graphical VNC client. If you want, you can also start this application from any terminal window by executing the `vinagre` command. This displays a window like the one shown in Figure B1-7.

FIGURE B1-7

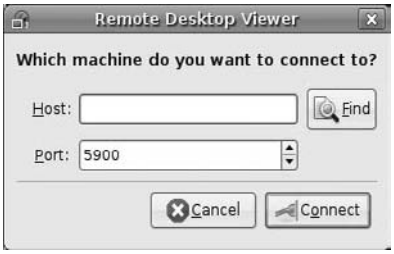
Ubuntu's Remote Desktop Viewer application, vinagre



Once this window displays, click the Connect icon to display a dialog like the one shown in Figure B1-8. Enter the name or IP address of the system that you want to connect to and (optionally) specify the number of the port on which its VNC server is running.

FIGURE B1-8

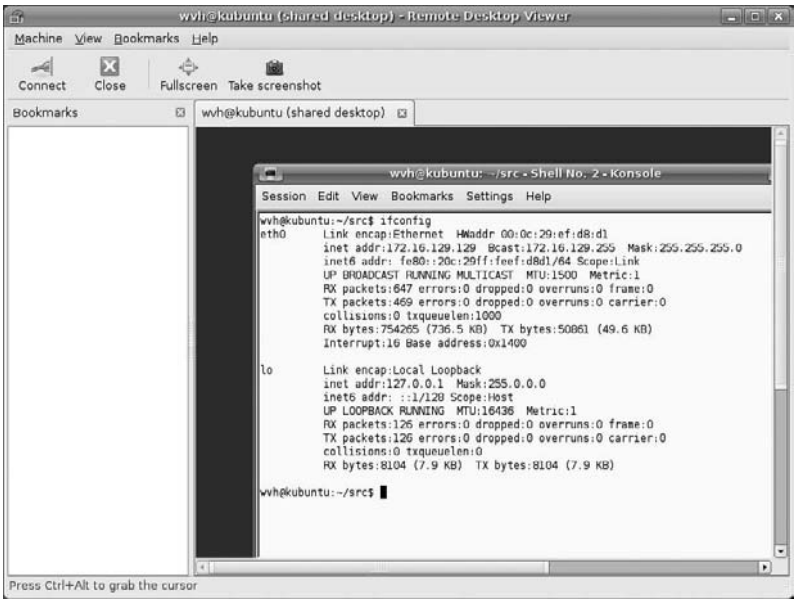
Identifying the system to connect to in vinagre



After entering this information, click the Connect button to connect to the VNC server on the specified system. If the remote system requires a password, a dialog displays in which you can enter that password. Once you have successfully connected to the remote system, a VNC connection to the specified system opens in a new tab in the window shown in Figure B1-7, as shown in Figure B1-9.

FIGURE B1-9

Viewing a Kubuntu system in the Remote Desktop Viewer



Two nice features of Ubuntu's Remote Desktop Viewer are its ability to open VNC connections to multiple systems at the same time, showing each in its own tab, and its ability to bookmark systems to which you've connected in order to simplify visiting them again — just as you can in a Web browser. To create a bookmark for a system that you've connected to, select the Bookmarks ➪ Add to bookmarks menu command, which enables you to provide your own name for the remote system. The new bookmark displays in the left-hand bookmarks pane. You can see examples of bookmarks in the next two figures.

Being able to connect to the graphical console of a remote system simplifies system administration and problem solving and can generally confuse and amaze your kids if they don't know that you're connecting to their machines. Figure B1-10 shows a VNC connection to a remote Mac OS X system using the Remote Desktop Viewer.

FIGURE B1-10

A VNC connection to a Mac OS X system

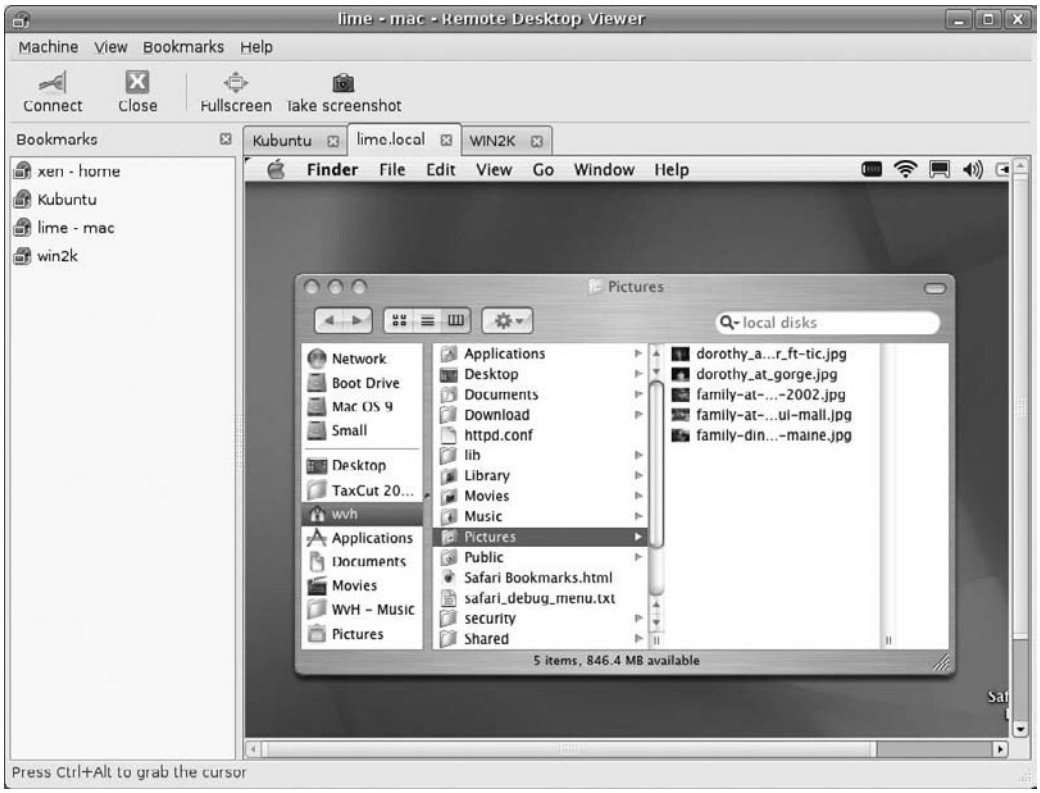
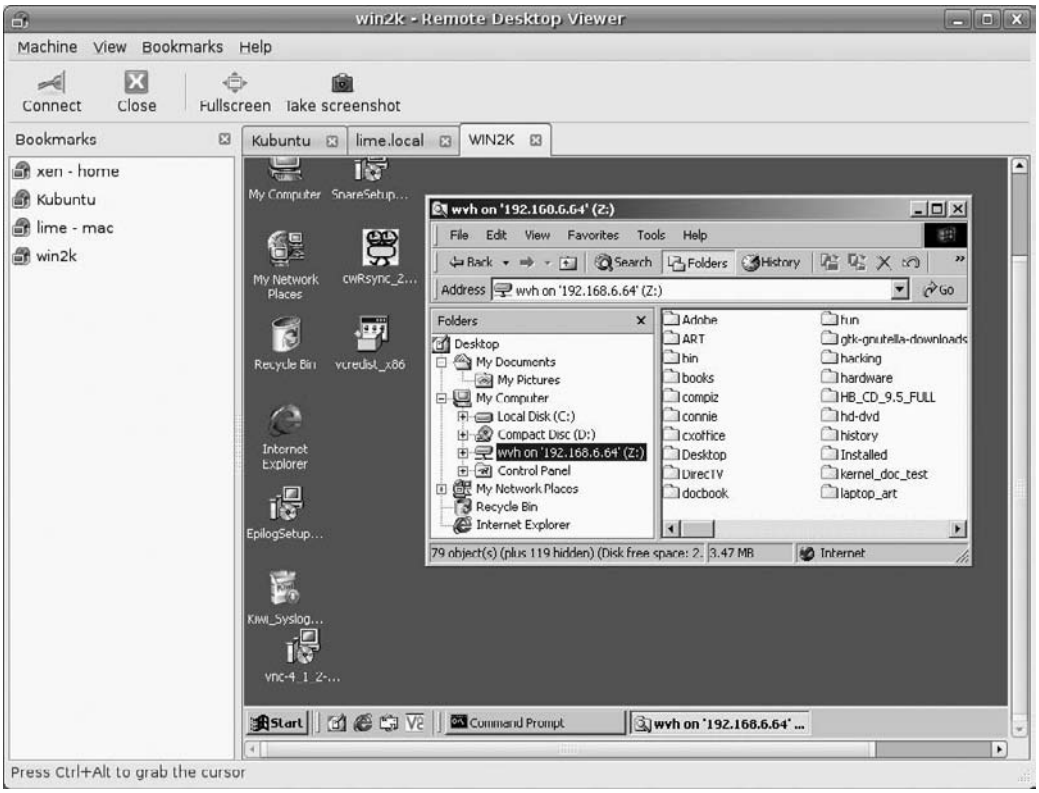


Figure B1-11 shows a VNC connection to a remote Microsoft Windows 2000 system using the Remote Desktop Viewer.

FIGURE B1-11

A VNC connection to a Windows 2000 system



Using Kubuntu's Remote Desktop Connection

Once a VNC server is running on one of your systems, you can connect to it from any Kubuntu system by selecting the K Menu ➪ Internet ➪ Krdc - Remote Desktop Connection menu command to start Kubuntu's graphical VNC client. If you want, you can also start this application from any terminal window by executing the `krdc` command. This displays a window like the one shown in Figure B1-12.

TIP **If the top part of this dialog does not show a browser list, click the Browse button to display it. Similarly, you can click the Browse button again to hide this portion of the dialog.**

Once this window displays, the Remote Desktop Connection application searches the network for available VNC servers and displays name and address information about each. You can either select the name of the system from the list at the top of the dialog (if it is shown), or enter the name or IP address of the host that you want to connect to in the Remote Desktop text area in forms such as `vnc:/`*`hostname`*, `vnc:/`*`IP-address`*, *`hostname:display-number`*, *`IP-address:display-number`*, and so on. Click the Connect icon to display a dialog like the one shown in Figure B1-13.

FIGURE B1-12

Kubuntu's Remote Desktop Connection application, krdc

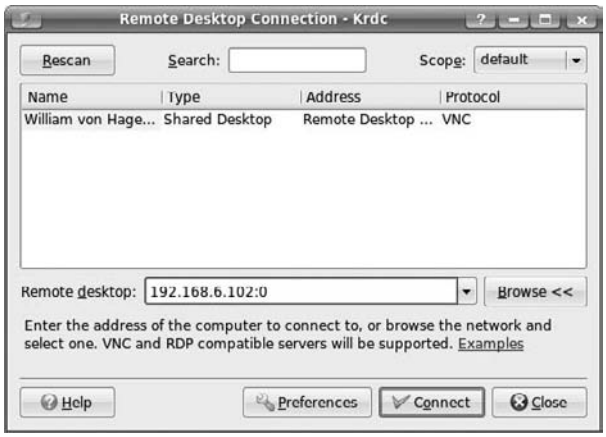
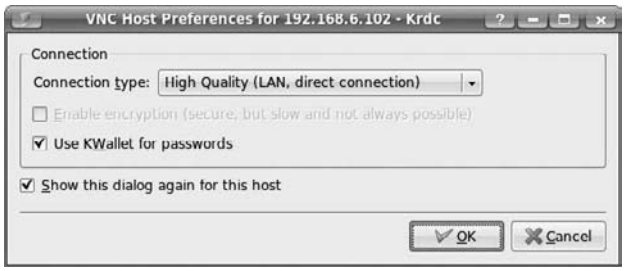


FIGURE B1-13

Specifying connection parameters for a remote system in krdc

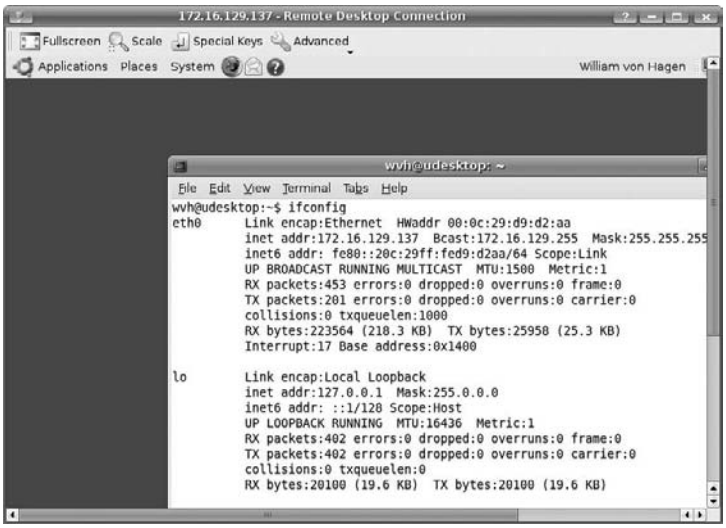


NOTE **Simply entering a hostname or IP address with no display number in the Remote Desktop text area will not work. You must either specify a display number (typically, :0), or use the** `vnc:/`*`hostname`* **or** `vnc:/`*`IP-address`* **specification to avoid having to do so. The Connect button will become active only after a valid VNC host specification is entered.**

This dialog displays the first time that you connect to any VNC server, and lets you set default values for this and subsequent connections to that server. The Connection type drop-down enables you to specify lower-quality VNC connections that will perform better on lesser-quality network connections. Selecting the Use KWallet for passwords checkbox specifies that the password to the remote system will be securely stored in the KWallet application's database. Once you have made any changes that you want to make, click OK to connect to the remote system. If you have elected to store your password in KWallet, you will be prompted for your password in order to open your wallet. If the remote system requires a password, a dialog displays in which you can enter that password. Once you have successfully connected to the remote system, a VNC connection to the specified system opens in a new dialog, as shown in Figure B1-14.

FIGURE B1-14

Viewing an Ubuntu system from the Kubuntu Remote Desktop Connection



Any connection to a remote system using krdc displays a toolbar across the top of the screen that provides some controls for that krdc window. These controls are as follows:

- **Fullscreen:** Increases the size of the krdc window to fill the entire screen. This can be especially useful when connecting to a system whose screen is larger than that of the one from which you are connecting, because it will display more of the screen area from the remote system.
- **Scale:** Reduces the size of the remote screen image so that it fits in the current krdc window. This can be useful when connecting to a system whose screen is larger than that of the one from which you are connecting, because it will display the entire screen from the remote system. On the down side, the image can be somewhat hard to read if the remote screen is very large.
- **Special Keys:** Enables you to send special key combinations to the remote system, which can be useful if those key combinations would also have an effect on your Kubuntu system.
- **Advanced:** Provides two configuration commands: View, which makes the remote desktop read-only (you can't change anything), and Always show local cursor, which can make it easier for you to track cursor movement on the remote system by displaying different cursors for the local and remote systems.

You can use the Remote Desktop Connection software to connect to multiple remote systems at the same time by starting multiple iterations of this application at the same time, one for each remote system that you want to connect to. Each connection to a remote system will have its own window.

Being able to connect to the graphical console of a remote system simplifies system administration and problem solving and can generally confuse and amaze your kids if they don't know that you're connecting to their machines. Figure B1-15 shows a VNC connection to a remote Mac OS X system using the Remote Desktop Connection.

Figure B1-16 shows a VNC connection to a remote Microsoft Windows 2000 system using the Remote Desktop Viewer. Note that these figures are essentially identical to those shown in Figure B1-10 and B1-11 except for the application that is being used.

FIGURE B1-15

A VNC connection to a Mac OS X system

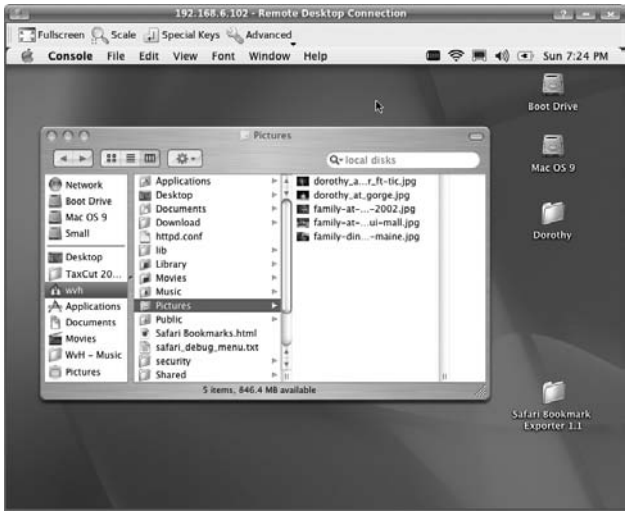


FIGURE B1-16

A VNC connection to a Windows 2000 system



Connecting to Remote Windows Terminal Servers

For most people, needing to interact with systems running Microsoft Windows is a fact of corporate, and occasionally personal, life. In the home networking environment, most Windows connectivity can be handled quite nicely using the VNC server and client discussed earlier in this chapter, but if you need to connect to a remote Windows system at work or your home computing budget is a bit higher than mine, you may need to access Windows systems that are running a Windows Terminal Server.

If you need to give users occasional, personal access to a Windows system but want to minimize costs and administrative hassles rather than install one machine per user, a good solution is to install Windows Terminal Services on a reasonably robust Windows system, and purchase a small pool of Client Access Licenses that are distributed by the server, as needed. Remote clients can then attach to the Terminal Services server and run virtual Windows sessions in a window on their desktops. You then install on the terminal server or in shares defined in your user profiles the software that people need to use, and any remote user connected to the server can access it to run the necessary software. Fortunately, remote access to Windows Terminal Services doesn't even require a Windows system anymore. Ubuntu users can easily access Windows Terminal Servers using `rdesktop` or the Terminal Server Client (`tsclient`), two open source software packages that speak the Remote Desktop Protocol (RDP) used by Windows Terminal Services. Kubuntu users can access Windows Terminal Servers using `rdesktop` or the KDE Remote Desktop Connection (`krdc`) software that was discussed previously.

TIP **If you centralize Windows services by running Terminal Services on a Microsoft Windows Domain Controller or server, you'll need to make sure that the users who want to connect to it have the "Log in Locally" user right or belong to a group with that right. Otherwise, users will receive the message "The local policy of this system does not permit you to log on interactively" and you may still have to give everyone a Windows box.**

Using rdesktop for Terminal Server Connections

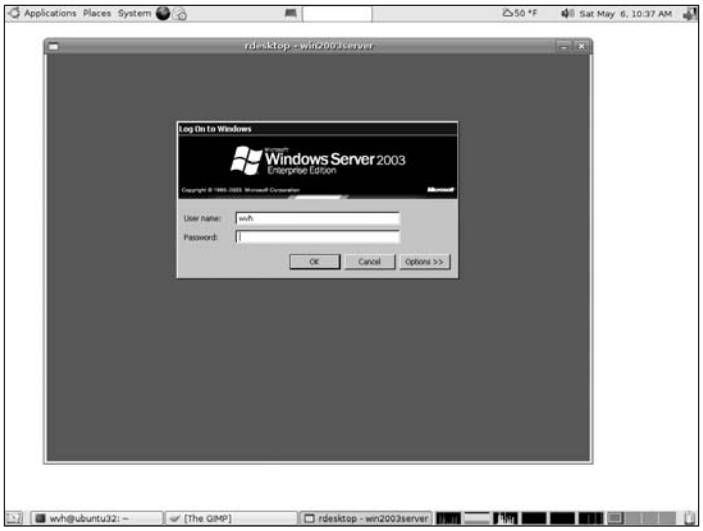
The `rdesktop` package (`www.rdesktop.org`) is a remote desktop package for Linux and most other UNIX and UNIX-like systems that run the X Window system. The `rdesktop` package is installed by default on both Ubuntu and Kubuntu Linux systems and provides an excellent, high-power client with a few really cool features as described in this section. The `rdesktop` package is available from the standard Ubuntu repositories, but you can always get the latest version of its source code and installable packages for other types of systems from its SourceForge site at `http://sourceforge.net/projects/rdesktop`.

You must start `rdesktop` from a GNOME or xterm terminal or from any KDE Konsole window, although you can certainly add it to your menus as explained in Chapters 5 and 6. The simplest `rdesktop` command line that you can use to connect to a remote system running Windows Terminal Services is `rdesktop` *`host`*, where *`host`* is the name or IP address of the system running Windows Terminal Services. A window displays on your desktop, showing the standard Windows login screen as shown in Figure B1-17.

After you specify the domain that you want to log in to (if necessary) and log in, your `rdesktop` window displays the standard Windows desktop, as shown in Figure B1-18, which was captured on a Kubuntu system to demonstrate that rdesktop works just fine on both Ubuntu and Kubuntu systems.

FIGURE B1-17

A Windows Terminal Services login prompt in rdesktop



The `rdesktop` application provides you with several options that can simplify access to Windows Terminal Services. Some of my favorites are the following:

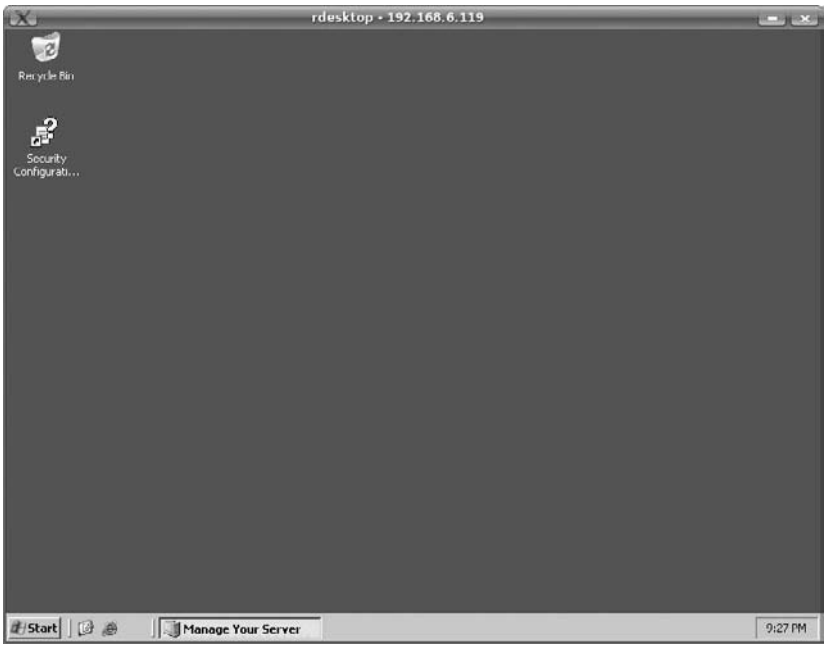
- `-d`: The domain that you want to authenticate to.
- `-f`: Full-screen mode. This displays `rdesktop` in a decoration-free window that takes over your desktop. You can toggle decorations (and therefore window controls) by pressing Ctrl+Alt+Enter.
- `-p`: Your password in the remote domain.
- `-u`: The name of the user that you want to log in as.

If the system running Windows Terminal Services is running Windows XP or a newer version of Windows, one especially handy option not listed in the previous section is the `-r` option, which lets you directly map resources on your Linux system to your Windows Terminal Services connection. This is especially useful when you want to map a local print queue to a virtual printer in your Windows Terminal session, or when you want to access a local drive in your terminal session. The syntax for these is `-r printer:local-queue-name` and `-r disk:share-name=/device/path`, respectively. For example, to attach PRN1 to a local print queue named SuperScript-660i, you would add `-r printer:SuperScript-660i` to your command-line options when executing the `rdesktop` command.

Similarly, to map your local CD-ROM drive to a share called cdrom, you would add `-r disk:cdrom=/dev/cdrom` to the `rdesktop` command line. If you still use floppies, you could map your local floppy drive to a share called floppy by adding `-r disk:floppy=/dev/fd0` to your `rdesktop` command line.

FIGURE B1-18

A Windows 2003 Server desktop in rdesktop



NOTE **The name that you specify as the share must consist of eight characters or fewer.**

Using Ubuntu's tsclient for Terminal Server Connections

In case you don't like `rdesktop` for some reason, Ubuntu also provides another Windows Terminal Server client as part of a default Ubuntu installation. This is the `tsclient` (Terminal Services Client) application (`http://gnomepro.com/tsclient`). The `tsclient` application is a GNOME-centric Terminal Server client that provides some optimization and also enables you to save information about a terminal server connection in a file that you can quickly and easily reload instead of manually specifying all of your settings each time that you want to connect. Although the latest version of `tsclient` for Ubuntu is always available from the Ubuntu repositories, you can get the latest source code and packages for other GNOME-based systems at the `tsclient` SourceForge project site at `http://sourceforge.net/projects/ts-client`.

To start the `tsclient` application, select Applications ➪ Internet ➪ Terminal Server Client, which displays the dialog shown in Figure B1-19.

If you enter only the name of the terminal server that you want to connect to and click Connect (or press Enter on your keyboard), the standard Windows Terminal Server login screen displays in a new `tsclient` window, and you can log in normally. Once you log in, a standard Windows desktop displays, as shown in Figure B1-20.

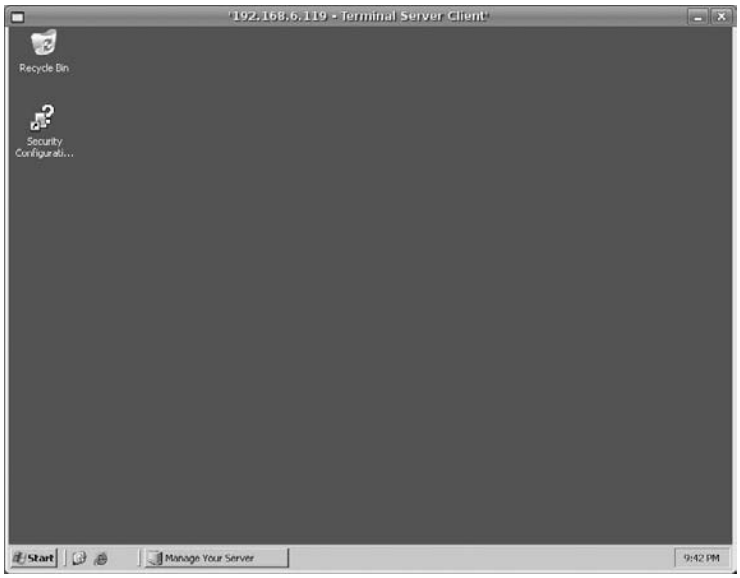
FIGURE B1-19

The tsclient startup screen



FIGURE B1-20

A Microsoft Windows terminal server session shown in tsclient



You can also customize your connection to a Windows Terminal Server by setting options that are available on tabs in the `tsclient` dialog. For example, the Local Resources tab, shown in Figure B1-21, enables you to make resources on your Ubuntu system available to your Windows Terminal Server session.

Similarly, the Display tab, shown in Figure B1-22, lets you specify the size of the window used to connect to the terminal server, which can be very handy when connecting to a terminal server from a system with a smaller screen, such as a laptop.

To simplify future connections to terminal servers that you may contact frequently, you can fill out all of the particulars for a connection in the dialog shown in Figure B1-19, click Save As, and enter the name that you want to associate with this connection. The `tsclient` application will save all of this information (except for your password) in a file by that name with the `.rdp` extension, and store it in the directory `~/.tsclient` in your home directory. In the future, you can select one of these files after starting `tsclient`, enter your password, and quickly establish a connection to the indicated terminal server.

Using Kubuntu's krdc for Terminal Server Connections

In addition to allowing you to use `rdesktop` to connect to Windows Terminal Servers from a Kubuntu system, the KDE Remote Desktop Connection (`krdc`) software that was discussed previously also supports RDP connections to Windows terminal servers. It does this by using the `rdesktop` software, which must therefore be installed on your Kubuntu system — because `rdesktop` is installed by default on both Ubuntu and Kubuntu systems, this works by default.

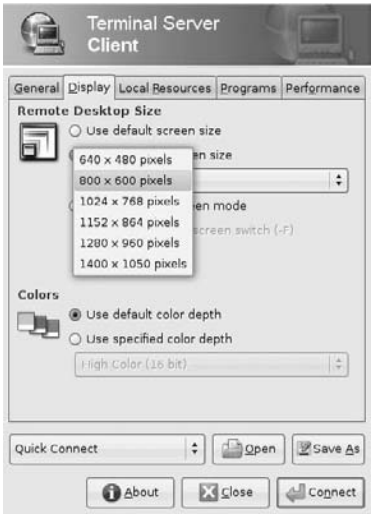
FIGURE B1-21

Making local resources available in tsclient



FIGURE B1-22

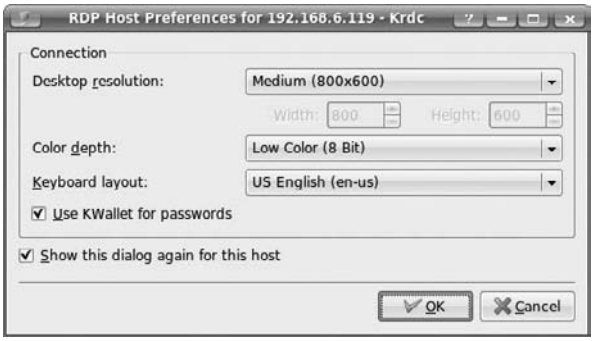
Specifying screen size in tsclient



Connecting to a Windows Terminal Server using krdc is remarkably similar to connecting to a VNC server using krdc. After selecting the K Menu ➪ Internet ➪ Krdc - Remote Desktop Connection menu command or executing the krdc command in a Konsole or xterm window, the screen shown in Figure B1-12 displays. Enter the name of the terminal server that you want to connect to in the Remote Desktop text area in the form `rdp://`*`hostname`* or `rdp://`*`IP-address`*, and press Return. This displays a slightly different dialog than the one shown in Figure B1-13. This dialog provides options that are specific to RDP connections, as shown in Figure B1-23.

FIGURE B1-23

Specifying RDP host preferences in krdc

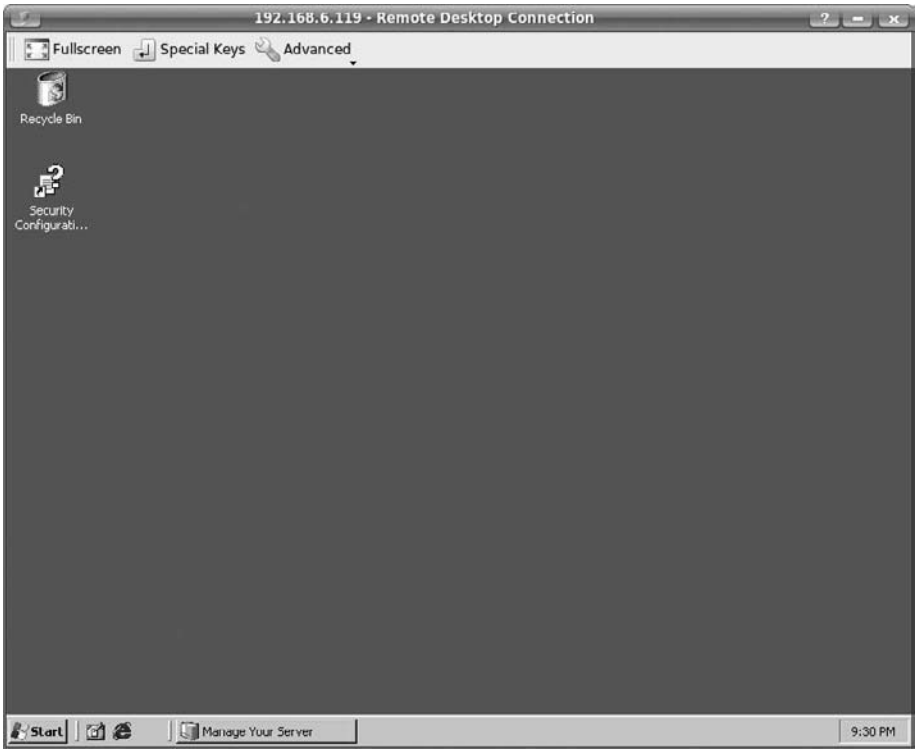


This dialog enables you to specify the size of the window used to connect to the terminal server, which can be very handy when connecting to a terminal server from a system with a smaller screen, such as a laptop. You can also specify the color depth to use — reducing the color depth can improve performance, especially on slower connections.

Once you have changed any settings that you want to update, click OK to proceed. If the KWallet option is still selected, you will be prompted for your Kubuntu password, after which a standard Windows Terminal Server login screen displays. After logging in, you will see a standard desktop, as shown in Figure B1-24.

FIGURE B1-24

A Windows Terminal Server connection in krdc



If you compare the screens shown in Figure B1-18 and Figure B1-24, it's easy to see that krdc is indeed using `rdesktop` to connect to the remote system. The only difference between the screens shown in these two figures is the presence of the standard krdc controls at the top of Figure B1-24. Note that the toolbar at the top of the screen shown in Figure B1-24 does not display the Scale command described in the earlier section of this chapter entitled "Using Kubuntu's Remote Desktop Connection." This is because the dialog shown in Figure B1-23 enables you to explicitly set the size of the remote desktop, and scaling is therefore not necessary once an RDP connection has been established.

Using Other Remote Connection Software

Although it is not supplied as part of a default Ubuntu or Kubuntu Linux installation, another popular software package for graphically connecting to remote systems is the free version of NoMachine NX, known as FreeNX. This package (named `freenx`) is not provided in the Ubuntu repositories and is not installed by default on a standard Ubuntu system, primarily because its Linux clients are KDE-based, and therefore install many KDE libraries. However, versions of FreeNX are available from some of the additional, user-supported repositories listed in Chapter 23, "Adding, Removing, and Updating Software." See that chapter for additional information.

FreeNX is a free, GPL implementation of NoMachine's NX (NoMachine NX) Server (`www.nomachine.com`). NoMachine has developed a compression technology that substantially reduces the size of X Window system communications and adds other performance improvements through caching and general protocol optimization. The NX client applications for Linux are included in the `freenx` package — the clients for Microsoft Windows, Mac OS X, and Solaris are available as free downloads from the NoMachine site. The free and commercial versions of the NX server differ primarily in terms of capabilities (and, of course, cost). FreeNX provides all of the core capabilities of the commercial NX server for remote connections, but does not currently include the SMB and printing (CUPS) support provided by the commercial NX server.

If you're already a VNC fan, NX is definitely worth a look for its superb compression technologies and associated performance improvements, and also because it inherently uses SSH for secure communications between client and server. FreeNX is effectively equivalent to VNC, and also supports Linux Terminal Server (`www.ltsp.org`), Windows Terminal Server, Citrix Metaframe, and standard X Window system XDMCP (X Display Manager Control Protocol) connections. However, it is not part of any official Ubuntu repository from Canonical, so you won't be able to get support from them if you encounter any problems. (You may still get free sympathy, however.)

Summary

Networking is a given in today's academic, business, and even home computing environments. The need to access more and different systems goes hand in hand with that. This chapter discussed secure command-line mechanisms for connecting to remote Linux and UNIX systems and explained how to install the SSH server on your system so that people can securely connect to your Ubuntu or Kubuntu system. I then explored the excellent VNC (Virtual Network Computing) clients and servers that enable you to remotely connect to the consoles of graphical systems such as Microsoft Windows, Mac OS X, Linux, and UNIX systems. The chapter concluded by discussing mechanisms for connecting to a Microsoft Windows Terminal Server from your Ubuntu or Kubuntu system, and then introduced the popular NX and FreeNX software that enables you to securely connect to just about any remote graphical console service.

Bonus Chapter 2 discusses the related topic of connecting to remote file servers to access files that are stored on them, and explains how to share files from your Ubuntu or Kubuntu system so that other systems on your network can quickly and easily access them. Music server, anyone?

Bonus Chapter 2

File Transfer and Sharing

Sooner or later, you're going to want to copy a file to or from another computer system. In pre-networking and pre-e-mail days (which you may not even remember), this usually involved writing a copy of the file to removable media of some sort, such as a floppy disk, carrying that floppy from one system to the other, and then copying the file from the floppy disk to the other system. This was known as "sneaker net"—you just can't beat that nerd humor!

Today, thanks to the fact that networking is almost ubiquitous and that most businesses and many homes have their own networks, electronic file transfer is an important capability. As you'd expect, there are many ways to transfer or copy files between systems, ranging from point-to-point file transfers to the peer-to-peer (P2P) file transfers much beloved of the music and film industries. In point-to-point file transfers, you explicitly copy a file from one system to another, and you therefore know the addresses or identities of both hosts. In peer-to-peer file transfer, groups of systems share files with each other without necessarily knowing or caring about the specific systems from which they are getting the files or to which they are sending them. In peer-to-peer file transfers, somebody obviously has to know what hosts are sharing what files, but this is usually built into the software—more about that later in this chapter.

The first section of this chapter discusses popular point-to-point file transfer solutions that are built into your Ubuntu or Kubuntu Linux system. The second section of this chapter, "Peer-to-Peer File Sharing on Ubuntu," discusses various P2P solutions that are easily installed and used on your system.

NOTE **You may have noticed by now that I am not a lawyer and that I have significant disdain for industry groups that try to suppress technologies just because they can be used to do something that may be illegal. Heck, I could conceivably club someone to death with my laptop, but does that mean that we should outlaw laptops? There is absolutely nothing inherently illegal about peer-to-peer file sharing. P2P is a great way to efficiently share and transfer large files among multiple systems, with everyone sharing part of the load. What you do with it is your business. Explaining how you can use it is mine.**

IN THIS CHAPTER

Transferring files between systems

Sharing files with Windows systems

Sharing files with Linux and UNIX systems

Peer-to-peer file sharing

Using BitTorrent

Transferring Files to Other Systems

As noted in the introduction to this chapter, this section discusses how to copy files between systems in cases where you know the identities of the machines involved in the transfer. Common examples of this sort of file transfer are "I need to copy this file from my laptop to the machine on my desk," "I want to put this file on our home media server, so we can play it from any machine in the house," or "I need to put this spreadsheet on the corporate file server so that my boss can see that I'm actually working from home." In each of these cases, you have authorized access to the machines at both ends of the transfer and you know their names, whether by fully qualified hostname (ftp.vonhagen.org), IP address (207.44.142.34), or their local hostname on your network (laptop, server1, and so on). You'll also know the name of the directory in which the file that you're looking for is located, and the name of the Windows share in which that file is located for SMB/CIFS file transfers.

In addition to having authorized access to the system that you're transferring files to, the appropriate service already must be running on the remote system, holding up a sort of electronic catcher's mitt that can receive the files that you're transferring there. This chapter focuses on the client-side of file transfers—transferring files to something that's ready to receive them, and doing anything you need to do to enable your local machine to reach them. This chapter doesn't discuss how to set up shared directories on Microsoft Windows or Linux systems. For information on sharing Linux directories so that you can reach them directly from a PC running Microsoft Windows, Mac OS X, or Linux, see Bonus Chapter 14, "Setting Up a Samba Server." For information on sharing Linux directories via NFS so that you can reach them directly from a PC running Mac OS X or Linux, see Bonus Chapter 13, "Setting Up an NFS Server." (You can also contact NFS Servers from Microsoft Windows but only by installing special software that doesn't come with Windows by default.) Information about setting up an FTP server is provided as a sidebar later in this chapter, because I received a few requests for information on that topic from readers of the first edition of this book.

TIP **Another common file transfer scenario is "my brother would really like a copy of this file." In that case, I'd suggest just attaching the file to an e-mail message and sending it to him, as long as the file in question is under the size of any quota that is imposed by the recipient's e-mail server. You don't know the name of his machine or its IP address, and you probably don't have direct access to the machine anyway.**

Creating Permanent Desktop Icons for Services and Locations

The traditional way to connect to remote services in the graphical GNOME and KDE desktop environments is to use the file manager for those systems: Nautilus in GNOME and Konqueror or Dolphin in KDE. Using Nautilus to connect to FTP, SSH (Secure FTP), SMB/CIFS, and other servers creates a temporary icon for those servers that is automatically removed when you log out. You can also bookmark these servers, which is permanent but only visible from the Places menu or within Nautilus. Similarly, Konqueror and Dolphin create permanent icons for remote servers that you connect to, but they are only visible within the network sections of those file managers.

The easiest way to create a permanent desktop icon on either a GNOME or a Kubuntu system is to create a desktop shortcut to the remote location:

- On GNOME systems, you can create a Launcher for a remote location by right-clicking on the desktop, selecting Launcher from the context-sensitive menu, and selecting Location as the Type of that launcher. The dialog will now look like the one shown in Figure B2-1.

FIGURE B2-1

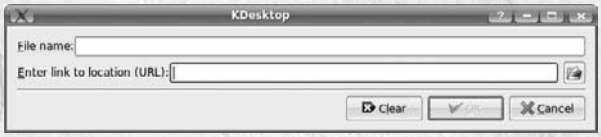
Creating a desktop launcher for a location in GNOME



- Enter some memorable text as the name of the location and identify the remote server by entering a URL of the form `PROTOCOL://USER@HOST/SHARE/DIRECTORY`. Replace `PROTOCOL` with `ftp`, `smb`, or `ssh`; replace `USER` with the name of the user that you want to connect as (you can leave out the `USER@` portion of the URL if you are connecting to an anonymous FTP server); and specify the hostname or IP address of the system that you want to connect to. When connecting to a shared Windows directory, you must also specify the name of the `SHARE` that you want to connect to. Finally, you can specify a `DIRECTORY` path on the host or `SHARE` that you want to connect to, relative to the `SHARE`, login directory, or anonymous FTP directory.
- Click Save to save the new launcher. Double-clicking the desktop icon for that launcher will display an authentication dialog for the remote service if authentication is required. Entering the correct password will then display the remote location in Nautilus. (SSH connections are displayed as Secure FTP connections.)
- On KDE systems, you can create a link to a remote location by right-clicking on the desktop and selecting the Create New ➪ Link to Location (URL) menu command to display an empty version of the dialog shown in Figure B2-2.

FIGURE B2-2

Creating a desktop link for a URL in KDE



- Enter some memorable text as the name of the location and identify the remote server by entering a URL of the form `PROTOCOL://USER@HOST/SHARE/DIRECTORY`. Replace PROTOCOL with `ftp`, `smb`, or `ssh`; replace `USER` with the name of the user that you want to connect as (you can leave out the `USER@` portion of the URL if you are connecting to an anonymous FTP server); and specify the hostname or IP address of the system that you want to connect to. When connecting to a shared Windows directory, you must also specify the name of the `SHARE` that you want to connect to. Finally, you can specify a `DIRECTORY` path on the host or `SHARE` that you want to connect to, relative to the `SHARE`, login directory, or anonymous FTP directory.

continued

continued

- Click Save to save the new link. Double-clicking the desktop icon for that link will display an authentication dialog for the remote service if authentication is required. Entering the correct password will then display FTP and SMB connections in Dolphin, or SSH connections in a Konsole window.

If you're a fan of desktop icons for remote services, rather than selecting a menu item in Ubuntu or pre-opening a file manager on KDE systems, creating a desktop launcher/link is "the right thing to do" because these items persist after you log out and log back in again.

Using FTP

FTP in uppercase letters stands for File Transfer Protocol, and is the oldest existing way of transferring files over a network, using the standard TCP/IP protocol that is the basis of the Internet. In lowercase letters, ftp is the name of a command-line program that you can run to connect to a remote FTP server so that you can transfer a file in either direction. As I'll explain in this section, Ubuntu makes it easy to transfer files using FTP in a variety of ways: the GNOME and KDE desktops provide built-in support for FTP connections, Ubuntu comes with the command-line ftp client, and the Ubuntu repositories provide standalone, graphical FTP applications if neither of the previous solutions meets your needs.

TIP **If you need to access files from an FTP server and aren't interested in using a GNOME-specific, KDE-specific, or command-line FTP client approach, you can simply connect to the FTP server using your favorite Web browser. For example, if the files that you're looking for are on an FTP server on a host named `xen.vonhagen.org`, you can try connecting to that host from your browser by specifying a URL of the form ftp://*USERNAME*@xen.vonhagen.org, where `USERNAME` is the account that you want to use to connect to the FTP server. If the FTP server supports anonymous connections, you can connect anonymously using a URL of the form ftp://xen.vonhagen.org.**

FTP is an old protocol, and the commands that you need to know are simple, even from the command line. You may never need to use ftp from the command line at all, because many equivalent command-line tools exist, such as wget and curl, which use newer protocols such as HTTP to retrieve files by directly specifying their URLs. In general, I prefer to err on the side of caution in this book—if you don't need certain information, that's fine and you can skip over it. However, if you do need to use FTP and information about it wasn't present in the book, both of us would be unhappy.

NOTE **FTP is as old as TCP/IP networking. If you're wondering why I'm even bothering to talk about FTP, the reason is that it may be the only mechanism that you can use to transfer relatively large files. Many mail servers enforce a size limit of 10MB on attachments, which is a trivial size for archive files, video files, and so on. If you can't access a shared directory that contains the files that you need to download or upload and they aren't available for P2P download as a torrent, FTP is the tried and true solution that you'll probably need to use.**

Connecting to an FTP Server from the GNOME Desktop

Ubuntu's Places menu provides the `Connect to Server` command, which makes it easy to create an icon on your desktop for an FTP server, and then to connect to that share using the Nautilus file manager. Selecting Places ➪ Connect to Server displays the dialog shown in Figure B2-3.

FIGURE B2-3

Defining a connection to a Public FTP server



The default selection, Public FTP, is not the one that you want to use if the FTP server that you want to connect to requires authentication (a specific username and password), which is usually the case. (I'll explain how to do that a little bit later in this section.) As you might suspect, Public FTP servers are servers that do not require authentication and are traditionally known as *anonymous FTP servers*. However, if you are actually connecting to an anonymous FTP server, you must at least enter the name or IP address of the FTP server that you want to connect to. You need to specify a specific port only if the remote FTP server is running on a nonstandard TCP/IP port, and you need to specify the name of a remote directory only if you want to automatically connect to some directory other than the default one that the FTP server associates with the anonymous user. If you are indeed trying to connect to an anonymous FTP server, click Connect to create a desktop icon for that FTP server and to attempt to connect to it.

TIP **If you expect to connect to the same anonymous (or authenticated) FTP server multiple times, you can click the Add bookmark checkbox and enter the name of a bookmark that will be added to the Network section of your GNOME desktop's Places menu after you successfully connect to the FTP server for the first time.**

To specify that you want to connect to an FTP server that requires authentication, click the Service type menu at the top of this dialog and select the FTP (with login) entry to display the dialog shown in Figure B2-4.

As with an anonymous FTP server, you must at least enter the name or IP address of the FTP server that you want to connect to. You can also identify the name of the user that you want to connect as, but this is optional—if you don't specify one, you'll be prompted for that information when you actually connect. As with the Public FTP server, you need to specify a specific port only if the remote FTP server is running on a nonstandard TCP/IP port, and you need to specify the name of a remote directory only if you want to automatically connect to some directory other than the default one that the FTP server associates with any user that you've specified. Click Connect to create a desktop icon for that FTP server and to attempt to connect to it.

When connecting to an authenticating FTP server, you will see a dialog that prompts you for any required information that you have not specified yet, as shown in Figure B2-5.

FIGURE B2-4

Defining a connection to an FTP server that requires authentication



FIGURE B2-5

Completing the login information for an FTP server



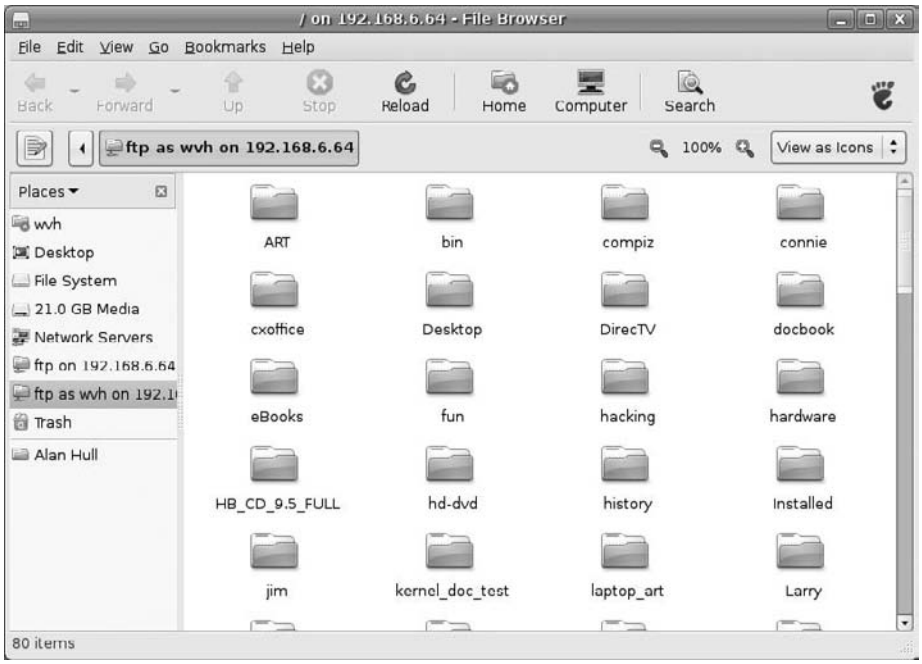
If you are connecting to an FTP server that requires authentication and didn't specify the name of a user to connect as, the dialog shown in Figure B2-5 still gives you the opportunity to connect anonymously. However, to actually log in, select the Connect as user radio button and enter the login name and password of the user that you want to connect as.

Once you enter the correct user name and login (or successfully connect to an anonymous FTP server), a Nautilus file manager window appears, which displays the contents of the default directory associated with the user that you are logged in as, as shown in Figure B2-6.

Once the Nautilus file manager appears, you can copy files in either direction using the standard Nautilus and GNOME desktop conventions, as discussed in Chapter 5.

FIGURE B2-6

A connection to an FTP server in the Nautilus file manager



Once you are done accessing the FTP server, close the Nautilus file manager window. To discard the desktop icon, right-click on the icon for the FTP server and select Unmount Volume from the pop-up menu to sever the connection and discard the icon. The icons for FTP servers are transient and are automatically removed when you log out. (See the sidebar entitled "Creating Permanent Desktop Icons for Services and Locations" for information on how to create a permanent desktop icon if you want one.) To preserve a shortcut to an FTP server in GNOME, create a bookmark for it when you initiate the connection. This bookmark will be displayed in the Network section of the Places menu.

TIP **To remove a bookmark for an FTP server that you no longer plan to connect to, select the Places ⇨ Network menu item from the GNOME desktop. This displays your network bookmarks in Nautilus. You can then select the Bookmarks ⇨ Edit Bookmarks menu command; select any bookmark and click Remove to delete it. You can also do this from within any existing Nautilus window.**

Connecting to an FTP Server from the KDE Desktop

Kubuntu provides a number of ways of graphically connecting to an FTP server thanks to the richness of the KDE desktop and the power of file managers such as Konqueror and Dolphin. These mechanisms are handiest if you're going to be connecting to the same FTP server multiple times. For one-shot FTP connections, you can always use the command-line mechanism described in the next section, but why not take advantage of a desktop environment if you're going to use one in the first place?

The traditional way to graphically connect to an FTP server and create an icon for that connection is to do so within the Konqueror file manager (and, nowadays, in the Dolphin file manager). These approaches create an icon for the connection that you can see in the remote:/ section of Konqueror or Dolphin, but which isn't present on the standard KDE desktop. If you want to create a desktop icon for this FTP server, see the sidebar entitled "Creating Permanent Desktop Icons for Services and Locations," earlier in this chapter, for more information.

The next two sections explain how to connect to an FTP site in your KDE file manager of choice.

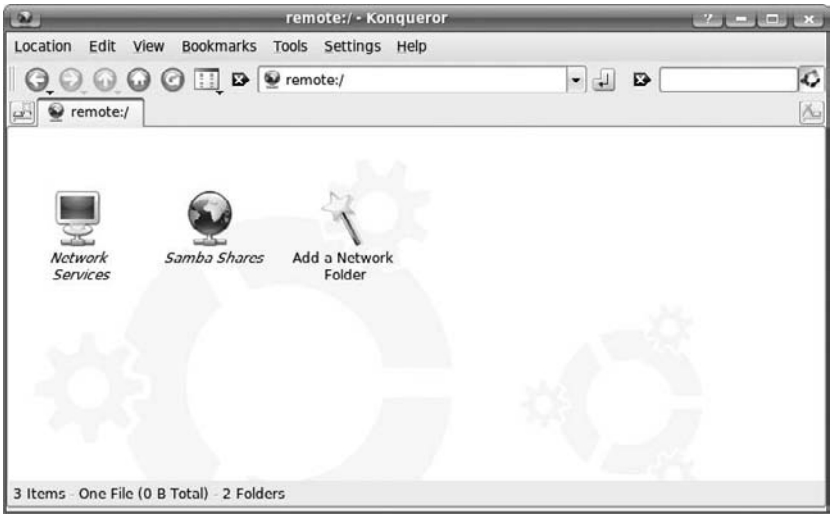
Connecting to an FTP Server in Konqueror

The traditional mechanism for connecting to an FTP server and creating a bookmark/icon for that connection is Konqueror, KDE's object browser for storage devices, network resources, Web pages, the kitchen sink, and so on. To create a connection to an FTP server, start Konqueror and open the remote:/ URL in the current or a new tab. This displays the page shown in Figure B2-7.

TIP **If you haven't changed the default location that is displayed when you start Konqueror, you can go to the remote:/ URL by clicking the Network Folders link that is displayed on its default "home page."**

FIGURE B2-7

Konqueror's default network folders page



To create a connection to an FTP server, click the Add a Network Folder icon. This displays the initial dialog for KDE's Network Folder Wizard, shown in Figure B2-8.

Select the FTP radio button, and click Next to proceed to a second dialog in which you can specify the details of the FTP server that you want to connect to, as shown in Figure B2-9.

FIGURE B2-8

The initial dialog for the KDE Network Folder Wizard



FIGURE B2-9

Defining an FTP connection in the KDE Network Folder Wizard



Enter a memorable name for the FTP server that you want to connect to in the Name field, and enter the name or IP address of the FTP server in the Server field. The Port field enables you to specify a particular port to connect to if the FTP server is running on a port other than the default port, port 21. If you always want to connect as a specific user, enter the name of that user in the User field; otherwise, leave that field blank. If you want to connect to a specific directory other than the default login directory associated with

whatever user you connect as, enter the full pathname of that directory in the Folder field. The Create an icon for this remote folder checkbox is selected by default, but you can disable this if you only want to connect to an FTP server once or simply don't want to clutter Konqueror's remote:/ page with icons.

After entering all of the relevant information for the FTP server that you want to connect to, click Save & Connect to connect to the FTP server and create a Konqueror icon for that connection if you left the Create an icon . . . checkbox selected. Konqueror displays the authentication dialog shown in Figure B2-10. If you did not specify a User in the dialog shown in Figure B2-9, you will have to enter the name of the user that you want to connect as; otherwise, you will simply have to enter the specified user's password on the FTP server.

FIGURE B2-10

An authentication request in KDE



TIP **To connect to an anonymous FTP server where you do not have a specific login, enter the username anonymous. It is traditional, but not mandatory, to provide your e-mail address as the password when connecting to an anonymous FTP server.**

Once you've successfully connected to an FTP server via Konqueror, Konqueror displays the contents of the default login directory or any login path that you specified, as shown in Figure B2-11. At this point, browsing the contents of the remote FTP site is the same as browsing a local directory. You can retrieve files from the remote site by dragging and dropping them to your desktop or another tab in Konqueror. Similarly, you can upload files (if you have permission to do so) by dragging and dropping them from your desktop, another tab in Konqueror, or a directory listing in a Dolphin window or pane.

Figure B2-12 shows Konqueror's remote:/ page after creating an icon for a connection to an anonymous FTP server, ironically named anonymous.

Connecting to an FTP Server in Dolphin

Dolphin is KDE's latest file manager of choice, and supports the same capabilities for displaying and browsing remote network locations that Konqueror does. To initiate a connection to an FTP server and optionally create an icon for that connection, start Dolphin and select Network from the Bookmarks list in the leftmost pane. This displays a Network location pane like the one shown in Figure B2-13.

FIGURE B2-11

An FTP site in Konqueror

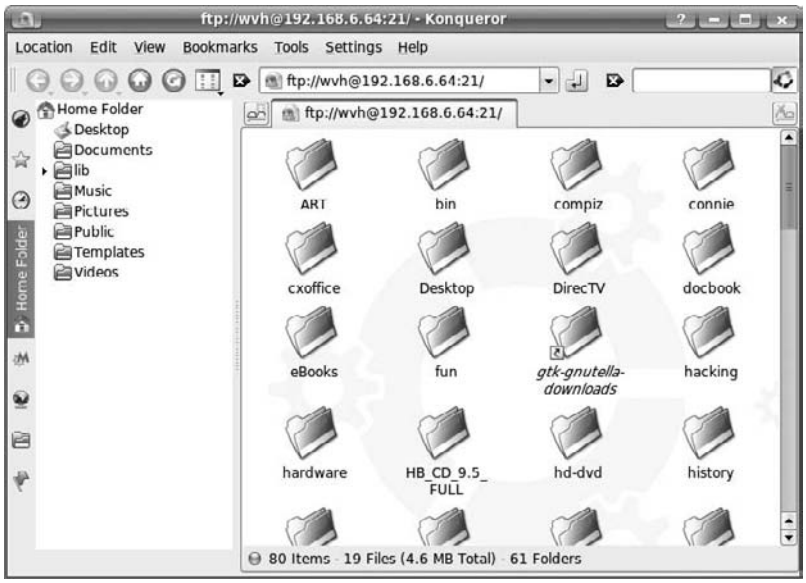


FIGURE B2-12

The icon for an FTP site named "anonymous" in Konqueror

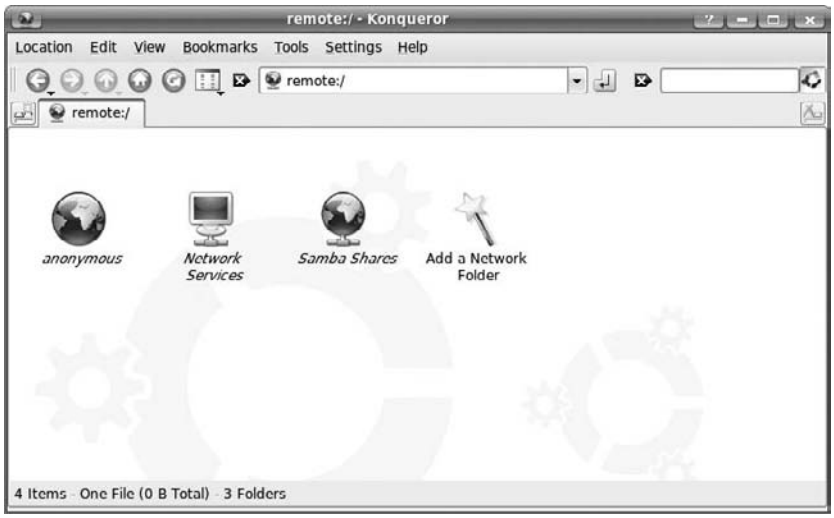
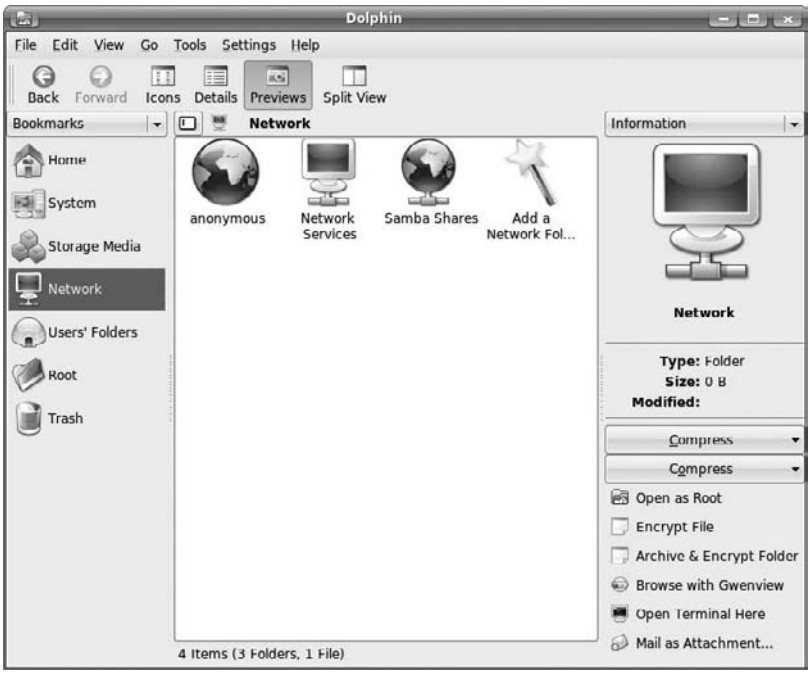


FIGURE B2-13

Dolphin's Network pane



You'll immediately notice that the contents of Dolphin's Network pane are identical to those of the remote:/ URL shown in Konqueror, down to the Add a Network Folder link to KDE's Network Folder Wizard. To use this icon to connect to an FTP server in Dolphin and optionally create an icon for that connection, follow the same steps that were described in the previous section, "Connecting to an FTP Server in Konqueror."

Connecting to an FTP Server from the Command Line

You can also connect to a remote FTP server by typing the `ftp` command and the name or IP address of the server from any Ubuntu or Kubuntu command line, such as the GNOME Terminal, KDE Konsole, or an xterm. The following example connects to an FTP server running on a system whose IP address is 192.168.6.64:

```
$ ftp 192.168.6.64
Connected to 192.168.6.64.
220-Welcome to Pure-FTPd.
220-You are user number 2 of 10 allowed.
220-This is a private system - No anonymous login
220-IPv6 connections are also welcome on this server.
220 You will be disconnected after 15 minutes of inactivity.
Name (192.168.6.64:wvh):
```

When you connect to an FTP server from the command line, you see some information about the FTP server and the system that it is running on, and the remote system eventually prompts you for the login that you want to use to connect to it. (How much information an FTP server displays about itself depends on the specific FTP server software that it is running and how the FTP server was configured.)

As you'll see from the last line of the previous example, an FTP server's login prompt usually tries to provide some defaults. In this case, I was logged in as the user "wvh" on the system from which I initiated the FTP connection. If that is the login name that you want to use to connect to the remote FTP server, you can simply press return to accept the default. If you want to connect as another user, type that user's login name and press return, as in the following example where I specify the user name "vonhagen":

```
Name (192.168.6.64:wvh): vonhagen
331 User vonhagen OK. Password required
Password:
```

TIP **You can usually log in on an anonymous FTP server from the command line by entering the login name "anonymous" and supplying your e-mail address as the password.**

After entering your login name, the remote FTP server prompts you for the password for the specified user on the FTP server. Enter that password and press return. If you've entered a valid username and password, you'll see a confirmation message, some optional information about the remote FTP server, and the `ftp>` prompt from the remote machine, telling you that it's ready to do something, as in the following example:

```
230 Login successful
Remote system type is UNIX.
Using binary mode to transfer files.
ftp>
```

All FTP servers support numerous commands, which you can see at any time by typing a question mark and pressing return, as in the following example:

```
ftp> ?
Commands may be abbreviated.  Commands are:
!         debug       mdir      qc          send
$         dir         mget      sendport    site
account   disconnect  mkdir     put         size
append    exit        mls       pwd         status
ascii     form        mode      quit        struct
bell      get         modtime   quote       system
binary    glob        mput      recv        sunique
bye       hash        newer     reget       tenex
case      help        nmap      rstatus     tick
cd        idle        nlist     rhelp       trace
cdup      image       ntrans    rename      type
chmod     lcd         open      reset       user
close     ls          prompt    restart     umask
cr        macdef      passive   rmdir       verbose
delete    mdelete     proxy     runique     ?
ftp>
```

Because the chances are good that you'll never use most of these (and you can get a complete explanation of them from the online man page for the `ftp` command by using the `man ftp` command from any command line), I'll just discuss the most basic set of commands required to find and retrieve a file on a remote machine.

All FTP servers export a directory hierarchy, just like any directory on a Linux system. Therefore, you can use familiar commands like `pwd` to find out the name of the current directory, `cd` to change to a specified directory, and so on. If you've logged in to an FTP server as a specific user, chances are good that you'll initially find yourself in the system's standard file transfer directory (called `/pub`) or in that user's home directory.

Once you've navigated to the directory to which you want to transfer a file or from which you want to retrieve one, you issue either the `put` command (to send a file to the FTP server) or the `get` command (to retrieve a file from the FTP server). Before actually transferring a file, you should make sure that the file transfer will take place in the appropriate mode. Most FTP servers automatically set the transfer mode so that you can retrieve binary files, such as graphics, archive files (file formats such as `.zip`, `.tgz`, `.bz2`, `.tar`, `.pax`, and so on), music files, and so on. Just in case, I like to always make sure that the file transfer will take place in binary mode by issuing the `bin` command, as in the following example:

```
ftp> bin
200 TYPE is now 8-bit binary
ftp>
```

If you do not transfer files in binary mode, they will be transferred as ASCII files, which will corrupt any binary file that you transfer and prevent you from opening it correctly. If you transfer a file and then can't view, play, or dearchive it, you may have overlooked this step.

By default, when you send or retrieve a file, the FTP command doesn't display any status information until the transfer completes. Personally, I like a bit of feedback just to know something is actually happening. I, therefore, always use the FTP server's `hash` command, which causes the server to display a hash mark (#) for each kilobyte of data that it transfers. Typing the hash command once activates this feature; typing it a second time deactivates the feature. The following example shows hashing being activated before an FTP file transfer:

```
ftp> hash
Hash mark printing on (1024 bytes/hash mark).
ftp>
```

Now you're finally ready to actually send or retrieve a file. Sending a file to the remote machine is as easy as typing the `put` command followed by the name of the file that you want to transfer to the FTP server, as in the following example:

```
ftp> put JDSC-0.5.9a.zip
local: JDSC-0.5.9a.zip remote: JDSC-0.5.9a.zip
200 PORT command successful
150 Connecting to port 39972
################################################################################
################################################################################
################################################################################
################################################################################
################################################################################
################################################################################
################################################################################
####################
226-File successfully transferred
226 0.202 seconds (measured here), 2.36 Mbytes per second
499668 bytes sent in 0.19 secs (2526.9 kB/s)
ftp>
```

Retrieving a file from the FTP server is equally easy: Type the `get` command followed by the name of the file that you want to retrieve from the FTP server, as in the following example:

```
ftp> get JDSC-0.5.9a.zip
local: JDSC-0.5.9a.zip remote: JDSC-0.5.9a.zip
200 PORT command successful
150-Connecting to port 35431
150 488.0 kbytes to download
##################################################################################
##################################################################################
##################################################################################
##################################################################################
##################################################################################
##################################################################################
##################################################################################
####################
226-File successfully transferred
226 0.105 seconds (measured here), 4.53 Mbytes per second
499668 bytes received in 0.12 secs (4217.0 kB/s)
ftp>
```

TIP **When sending or retrieving a file, you can give it a different name on the FTP server or your local machine (respectively), by specifying the new file name after the name of the one that you want to send or retrieve.**

Once the file transfer is complete, you can terminate your connection to the FTP server and exit from the `ftp` command by typing the `quit` command, as in the following example:

```
ftp> quit
```

It doesn't get much simpler than that! The nice thing about learning to use the standard command-line `ftp` utility is that it is available on most systems (even from the Microsoft Windows command line), and the commands are essentially the same everywhere. Learn once, use many, I always say.

Connecting to an FTP Server Using Graphical Clients

If you need to use FTP but would prefer a standalone graphical FTP client application, two popular ones are available from the Ubuntu repositories:

- If you are using a standard GNOME-based Ubuntu system, you'll probably want to install `gtk-gftp`, a GNOME FTP client that uses the GTK graphical interface. Installing this package will create an entry on the Applications ➪ Internet menu called gFTP. This command makes it easy to graphically transfer files using FTP, and also enables you to bookmark FTP settings (hostname, username, password, and so on) just as you would do in your favorite Web browser. Figure B2-14 shows the `gtk-gftp` application displaying the contents of a remote FTP server (the right pane) and a local directory (in the left pane) to or from which you can drag and drop files to transfer them.
- On a GNOME system you can, of course, also install and use the FileZilla client that is discussed in the next bullet.
- If you are using a Kubuntu system, you can install the gFTP client, but doing so will also install all of the GNOME and GTK baggage that gFTP requires. Two popular graphical FTP clients that are native to KDE are KFTPgrabber (once known as Kftp, `www.kftp.org`) and KBear (`http://kbear.sourceforge.net`), neither of which is included in the Ubuntu repositories. However, the Ubuntu repositories do include a graphical non-GNOME FTP client known as FileZilla

(`http://filezilla-project.org`), which uses a platform-independent set of widgets and therefore provides a relatively lightweight graphical FTP client for Kubuntu systems (and also works great on Ubuntu systems). Figure B2-15 shows the FileZilla application displaying the contents of a remote FTP server (the right pane) and a local directory (in the left pane) to or from which you can drag and drop files to transfer them.

Neither of these graphical FTP clients is installed on Ubuntu or Kubuntu systems by default, but both of them are easy to install using any of the software installation mechanisms described in Chapter 23.

FIGURE B2-14

Using the gtk-gftp graphical FTP client

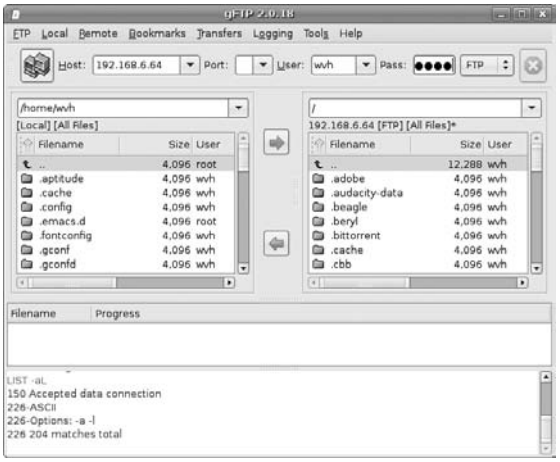
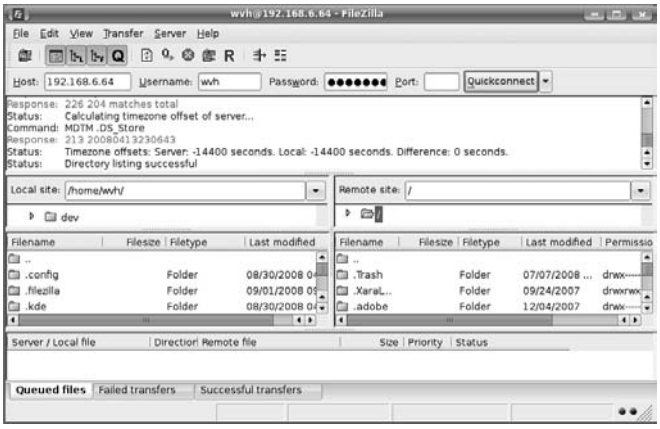


FIGURE B2-15

Using the FileZilla graphical FTP client



Setting Up an FTP Server

Many FTP servers are available in the Ubuntu repositories: `ftpd`, `inetutils-ftpd`, `krb5-ftpd`, `muddleftpd`, `proftpd`, `pure-ftpd`, `pyftpd`, `twoftpd`, `vsftpd`, `wu-ftpd`, and `wzdftpd`. The FTP protocol has been around for a long time, and people have had plenty of time to develop FTP servers! I'm not going to compare and contrast each of these, since they all have their own feature sets and fanclubs, and, frankly, because I don't think many people will care. Of these, my favorite is Pure-FTPd, which is available as the `pure-ftpd` package in the Ubuntu repositories. Why? It's easy to install, light-weight, trivial to configure, secure, fully internationalized, and runs in a number of modes, depending on the kind of FTP server that you want to run.

After installing the `pure-ftpd` package using any of the software installation mechanisms discussed in Chapter 23, Pure-FTPd is ready to use in its default configuration. Any user with an account on the system where you installed Pure-FTPd can now connect to their home directory and transfer files to/from there via FTP. By default, Pure-FTPd disables anonymous connections—to enable anonymous connections, modify the file `/etc/pure-ftpd/conf/NoAnonymous` to contain the word "no" (by default, it contains the word "yes," which means that anonymous users are not allowed). Next, create a user named `ftp` on your system whose home directory is the directory that you want anonymous users to be able to access. (See Bonus Chapter 1 for information on creating users.) Finally, restart the Pure-FTPd server using the command `/etc/init.d/pure-ftpd restart`. Voilà! Anonymous users are now supported.

Pure-FTPd has an amazing number of command-line options that control how it is started and all of its different modes of operations, all of which you can specify in the `/etc/init.d/pure-ftpd startup/shutdown` script. Pure-FTPd makes it easy to set up virtual FTP servers, IP-address-specific FTP servers, upload and download quotas, and much more. It also comes with a variety of administrative commands, my favorite of which is the `pure-ftpwho` command, which displays information about all current FTP users and what they're doing.

If you're more comfortable administering your FTP server graphically, you can install the PureAdmin application (the `pureadmin` package) on Ubuntu or Kubuntu systems. PureAdmin uses the GTK+2.x widgets, and thus won't install all of GNOME. If you're running KDE 3.x on your Kubuntu system, a KControl module named KcmPureFTPd (the `kcmpureftpd` package) is also available, but that hasn't been updated for a while and thus has some issues with the latest Pure-FTPd configuration files—I'd also avoid it on general principles, since KControl is not the system configuration framework of choice for the latest versions of KDE.

Accessing Shares on Windows Systems

Your Ubuntu or Kubuntu Linux systems can easily access shared directories on Microsoft Windows systems (commonly known as *shares*) thanks to one of the most popular and useful open source software packages ever created, Samba (`www.samba.org`). Samba was originally written by Andrew Tridgell who has received numerous awards and accolades for it. Tridgell still works on it today, though thousands of others have contributed to Samba, and it is now a team effort. See Bonus Chapter 14, "Setting Up a Samba Server," for a bit more background information on Samba.

Samba takes its name from the SMB (Server Message Block) protocol that is the original underlying protocol used for networked file sharing on Microsoft Windows systems. Tridgell created Samba in a massive

feat of reverse engineering the protocol and how it worked, because Microsoft's specifications for SMB weren't publicly available (big surprise). Today, thanks largely to Tridgell and Samba, Microsoft has opened up the SMB specification, which is now part of its larger CIFS (Common Internet File System) specification.

As mentioned in the introduction to this part of the chapter, you can also use Samba to share files from a Linux system so that specified directories look like Windows shares and Windows users can access them. This requires that you set up a Samba server, which is explained in Bonus Chapter 14. This section focuses on accessing a shared Windows directory from the desktop of your Ubuntu Linux system and copying files to or from that Windows share.

Locating Windows Shares

If you're unsure of the shared directories and associated servers that are available on your local network, you can quickly list them all using the smbtree command, which will prompt you for your Windows password, as in the following example:

```
$ smbtree
Password:
WVH
    \\YM901252              YM901252
       \\YM901252\ADMIN$       IPC Service (YM901252)
       \\YM901252\IPC$         IPC Service (YM901252)
       \\YM901252\disk1
    \\HD-HTGLOB5
       \\HD-HTGLOB5\info       TeraStation utilities
        \\HD-HTGLOB5\share     TeraStation folder
        \\HD-HTGLOB5\capture   TeraStation folder
        \\HD-HTGLOB5\IPC$      IPC Service (HD-HTGLOB5)
        \\HD-HTGLOB5\ADMIN$    IPC Service (HD-HTGLOB5)
    \\XEN                   Samba Ubuntu Server
       \\XEN\wvh               Home Directories
       \\XEN\brother           Brother Hl-2070N - hl1250 Driver
       \\XEN\Silentwriter      Office Silentwriter
       \\XEN\Stylus_200        Epson Stylus Photo R200
       \\XEN\IPC$              IPC Service (Samba Ubuntu Server)
    \\UBUNTU                Desktop Ubuntu
       \\UBUNTU\wvh            Home Directories
       \\UBUNTU\PDF            PDF
       \\UBUNTU\IPC$           IPC Service (Desktop Ubuntu)
       \\UBUNTU\print$         Printer Drivers
    \\WINSERVER             Windows 2000 LAN Mgr
       \\WINSERVER\IPC$        Remote IPC
       \\WINSERVER\print$      Printer Drivers
       \\WINSERVER\wvh         Home Directory
       \\WINSERVER\office      Shared Enterprise Data
```

If you're unsure of the shares that a Windows (or Samba) that a specific server exports, you can also list them by using the `smbclient` command, which will prompt you for your Windows password, as in the following example:

```
$ smbclient -L 192.168.6.2
Password:
Domain=[WINSERVER] OS=[Windows 5.1] Server=[Windows 2000 LAN Mgr]
```

```
        Sharename       Type      Comment
        ---------       ----      -------
        IPC$            IPC       Remote IPC
        print$          Disk      Printer Drivers
        wvh             Disk      Home Directory
        office          Disk      Shared Enterprise Data
Domain=[WINSERVER] OS=[Windows 5.1] Server=[Windows 2000 LAN Mgr]
        Server               Comment
        ---------            -------
        Workgroup            Master
        ---------            -------
                             HD-HTGLOB5
        WVH                  YM901252
```

You can use the information from either of these examples to identify the name of the share that you want to mount and the server from which that directory is being shared. You will need this information to mount that share on your Ubuntu or Kubuntu system using any of the mechanisms described in the next few sections.

Mounting Windows Shares at Boot Time

If you know that you will frequently need to access files and directories in a Windows share, the easiest way to make that share permanently available is to add it to the list of filesystems that your Ubuntu or Kubuntu system mounts when that system boots. The section of Chapter 4 entitled "Automatically Mounting Filesystems at Boot Time" explains the format of the text file that contains the list of filesystems that your system knows about and that can be mounted at boot time. This section explains the specific entries that are necessary when mounting a networked Windows shared directory at boot time.

If you know the name and associated server for the Windows share that you want to mount, you have all of the information that you need to mount that share on your Ubuntu or Kubuntu system. If not, you can use the information from either of the examples in the previous section to get the share and server name. To mount that share on your system each time that you start your system, add an entry for the Windows share that you want to mount to the file /etc/fstab. For example, to mount the office share shown in the previous example on the directory `/mnt/office` (which must already exist), I would create an `/etc/fstab` entry like the following:

```
//192.168.6.2   /mnt/office  cifs \
    auto,credentials=/etc/cifs.login,uid=0,gid=users  0  0
```

(This entry would appear on a single line in the /etc/fstab file—the backslash is only present for formatting purposes.)

The only unique portions of this entry are the mount options, which have the following meanings:

- auto: Automatically mount the filesystem when the system boots.
- credentials=/etc/cifs.login: Identifies the name of a file that contains the username and password to authenticate with when mounting the filesystem. Because the /etc/fstab file is usually publicly readable, you do not want to embed this information in that file for security reasons. You should put this information in a separate file and change that file so that it can only be read by root (sudo chmod 600 /etc/cifs.login). This file should contain two lines of the following form:

  ```
  username=USERNAME
  password=PASSWORD
  ```

- USERNAME is the name of the user that you want to use to mount the Windows share on the remote host, and PASSWORD is that user's password.
- uid=0: Use User ID 0 when mounting the filesystem.
- gid=users: All files on the Windows share will be assigned to the group users. This enables users of your systems that belong to this group (with their Windows login names) to write to files that they can access on the remote system.

Mounting a Windows share via the /etc/fstab file will make that share available to all users of your system, enforcing file access security through the users entry in the /etc/group file and Windows permission settings on the files and directories in the share, which is fine for multiuser server systems. For personal desktop Linux systems, you may prefer to require that users of each system manually mount the Windows shares that they themselves need, which will enable you to more precisely leverage their personal security settings on the Windows server. The next three sections discuss how users can mount Windows shares themselves on their Ubuntu and Kubuntu desktop systems.

Connecting to Windows Shares from the GNOME Desktop

Ubuntu's Places menu provides the Connect to Server command, which makes it easy to create an icon on your desktop for a Windows share and connect to that share using the Nautilus file manager. Selecting Places ➪ Connect to Server displays the dialog shown earlier in Figure B2-3. To specify that you want to connect to an FTP server that requires authentication, click the Service type menu at the top of this dialog and select the Windows share entry to display the dialog shown in Figure B2-16.

You must at least enter the name or IP address of the Windows server that hosts the share that you want to connect to. If you want to provide your own name for the desktop icon that this command will create, you can enter that in the Name to use for connection field, which can be handy because the default name of the icon is the name of the Windows server, and you may want to connect to multiple shares on the same server. You can also identify the name of the user that you want to connect as, the name of the Windows domain or workgroup that you want to connect as, and the share that you want to connect to, but all of this is optional—you'll be prompted for anything that you don't specify here when you attempt to connect. As you can see in Figure B2-16, I prefer to put as much information as possible here, but that's up to you. Click Connect to create a desktop icon for the Windows share that you have defined.

FIGURE B2-16

Identifying a Windows share to connect to



Simply creating a desktop icon for a Windows share doesn't actually establish the connection. To do that, double-click on the appropriate desktop icon to open it. You will see a dialog that prompts you for any required information that you have not specified yet, as shown in Figure B2-17.

Once you enter any remaining information about the share that you want to connect to and the server on which it resides, a Nautilus file manager window appears which displays the contents of the default directory associated with the user that you are logged in as, as shown in Figure B2-18.

FIGURE B2-17

Completing authentication information for a Windows share



FIGURE B2-18

A connection to a Windows share in the Nautilus file manager



Once the Nautilus file manager displays, copying files in either direction uses the standard Nautilus and GNOME desktop conventions, as discussed in Chapter 5.

Once you are done accessing the Windows share, close the Nautilus file manager window. To discard the desktop icon, right-click on the icon for the Windows share and select Unmount Volume from the pop-up menu to sever the connection and discard the icon. Desktop icons for SMB shares will be automatically removed when you log out. If you want to create a permanent desktop icon, see the sidebar entitled "Creating Permanent Desktop Icons for Services and Locations," earlier in this chapter, for more information.

Connecting to Windows Shares from Kubuntu Systems

Kubuntu provides a number of ways of graphically connecting to Windows shares, thanks to the richness of the KDE desktop and the power of file managers such as Konqueror and Dolphin. These mechanisms are handiest if you're going to be connecting to the same share multiple times. For one-shot connections to a Windows share, you can always use the command-line mechanism described in the section entitled "Connecting to Windows Shares from the Command Line," but why not take advantage of a desktop environment if you're going to use one in the first place?

The traditional way to graphically connect to a Windows share and create a desktop icon for that connection is to do so within the Konqueror file manager (and, nowadays, in the Dolphin file manager). These approaches create an icon for the connection that you can see in the remote:/ section of Konqueror or Dolphin, but which isn't present on the standard KDE desktop. If you want to create a desktop icon for a Windows share, see the sidebar entitled "Creating Permanent Desktop Icons for Services and Locations," earlier in this chapter, for more information.

The next two sections explain how to connect to a Windows share in your KDE file manager of choice.

Connecting to a Windows Share in Konqueror

The traditional mechanism for connecting to a Windows share and creating a bookmark/icon for that connection is through Konqueror, KDE's object browser for storage devices, network resources, Web pages, the kitchen sink, and so on. To create a connection to a Windows share, start Konqueror and open the remote:/ URL in the current or a new tab. This displays the page shown in Figure B2-7. From this pane, you can click the Samba Shares icon to display a screen that shows existing Windows shares that were automatically located by scanning your network, or you can manually create an icon for connecting to a share as described in the remainder of this section.



If you haven't changed the default location that is displayed when you start Konqueror, you can go to the remote:/ URL by clicking the Network Folders link that is displayed on its default "home page."

To create a connection to a Windows share, click the Add a Network Folder icon. This displays the initial dialog for KDE's Network Folder Wizard, shown in Figure B2-8.

Select the Microsoft Windows network drive radio button and click Next to proceed to a second dialog in which you can specify the details of the Windows share that you want to connect to, as shown in Figure B2-19.

Enter a memorable name for the Windows Share that you want to connect to in the Name field, and enter the name or IP address of the Windows server in the Server field. Enter the name of the share that you want to connect to in the Folder field. The Create an icon for this remote folder checkbox is selected by default, but you can disable this if you only want to connect to a Windows share once or simply don't want to clutter Konqueror's remote:/ page with icons.

FIGURE B2-19

Defining a connection to a Windows Share in the KDE Network Folder Wizard



After entering all of the relevant information for the Windows share that you want to connect to, click Save & Connect to connect to the FTP server and create a Konqueror icon for that connection if you left the Create an icon. . . checkbox selected. Konqueror displays the authentication dialog shown in Figure B2-20, in which you must supply the name of the user that you want to connect as, and their password.

FIGURE B2-20

An authentication request in KDE



Once you've successfully connected to a Windows share via Konqueror, Konqueror displays the contents of the default login directory or any login path that you specified, as shown in Figure B2-21. At this point, browsing the contents of the Windows share is the same as browsing a local directory. You can retrieve files from the remote site by dragging and dropping them to your desktop or another tab in Konqueror. Similarly, you can upload files by dragging and dropping them from your desktop, another tab in Konqueror, or a directory listing in a Dolphin window or pane.

FIGURE B2-21

An FTP site in Konqueror

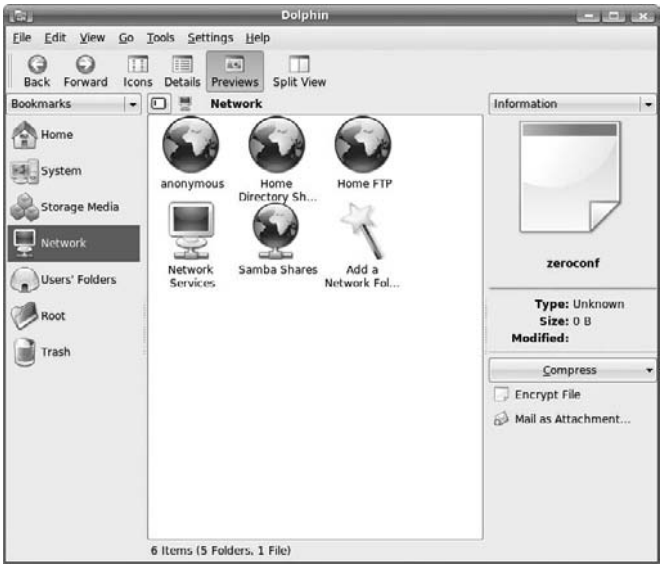


Connecting to an FTP Server in Dolphin

Dolphin is KDE's latest file manager of choice, and it supports the same capabilities for displaying and browsing remote network locations that Konqueror does. To initiate a connection to a Windows share and optionally create an icon for that connection, start Dolphin and select Network from the Bookmarks list in the leftmost pane. This displays a Network location pane like the one shown in Figure B2-22.

FIGURE B2-22

Dolphin's Network pane



You'll immediately notice that the contents of Dolphin's Network pane are identical to those of the remote:/ URL shown in Konqueror, down to the Add a Network Folder link to KDE's Network Folder Wizard. To use this icon to connect to an FTP server in Dolphin and optionally create an icon for that connection, follow the steps that were described in the previous section, "Connecting to an FTP Server in Konqueror."

Connecting to Windows Shares from the Command Line

While many people find using the command line to be a throwback to the 1980s and serial terminals, command-line commands are eminently portable, working fine from your terminal emulator of choice (a GNOME Terminal, KDE Konsole, xterm, and so on) on Ubuntu, Kubuntu, Xubuntu, Edubuntu, and even non-buntu Linux distributions.

To mount a Windows share from the command line on any Ubuntu or Kubuntu system, execute a command like the following:

```
sudo mount -t cifs //192.168.6.64/wvh /mnt/wvh \
    -o username=wvh,password=yeahright,rw
```

(This command would ordinarily be typed on a single line—the backslash in this command is only present for formatting purposes.)

This command mounts the Windows or Samba share `wvh` from the host 192.168.6.64 on the directory `/mnt/wvh` in read-write mode (`rw`), authenticating as the user wvh with the password yeahright. The mounted share will be available on this directory until it is unmounted, which you can do via the following command:

```
sudo umount /mnt/wvh
```

Note that this command isn't "unmount," but `umount`. Why type that extra "n" if you don't have to?

Accessing NFS Directories from Linux Systems

Long before Microsoft Windows, UNIX systems were sharing files over networks using Network File System (NFS), which was originally developed by Sun Microsystems, Inc. in the early 1980s. Sun was smart enough to release the specifications for NFS to the known universe, and every other UNIX system manufacturer with any common sense since then has also implemented support for NFS on their UNIX and UNIX-like systems. Today, most of those vendors are gone and even Sun's long-term future is cloudy, but NFS lives on as the default network filesystem and file sharing mechanism that you can find on any UNIX or UNIX-like system, such as Linux and Mac OS X.

Unfortunately, a whizzy graphical mechanism for mounting NFS filesystems isn't currently supported by the GNOME desktop's Connect to Server tool or KDE's Konqueror or Dolphin file managers, so you'll have to mount remote NFS filesystems manually, using the command line. And you may have thought you could get away with not reading Chapter 14! Luckily, mounting filesystems from the command line is extremely easy to do. (The mechanism by which you export a filesystem from your local system via NFS is similarly easy, and is explained in Bonus Chapter 13, "Setting Up an NFS Server.")

NFS filesystems are identified by the name or IP address of the NFS server on which they live, a colon, and the name of the directory that the server is exporting. For example, `192.228.6.64:/home/wvh` identifies the directory /home/wvh that is being exported by the NFS server whose IP address is 192.228.6.64. NFS filesystems are mounted on directories on your local Ubuntu or Kubuntu system just like any other filesystem (as explained in Chapter 4).

To mount the remote NFS filesystem that I just mentioned, you therefore need to create a mount point for it. I always create these in the directory `/mnt` so that they're easily identified as temporarily mounted devices. For example, the following command creates a mount point called `/mnt/nfs`:

```
$ sudo mkdir /mnt/nfs
Password:
```

After providing your password, the mount point exists, so you can now mount the NFS directory using the standard Linux `mount` command, as in the following example:

```
$ sudo mount -t nfs 192.228.6.64:/home/wvh /mnt/nfs
```

At this point, you can verify that the directory is successfully mounted either by listing that directory or by simply using the `df` command, as in the following example:

```
$ df
Filesystem           1K-blocks      Used Available Use% Mounted on
/dev/hda1             77299808  40359464  33013660  56% /
varrun                  257968       136    257832   1% /var/run
varlock                 257968         4    257964   1% /var/lock
udev                    257968        96    257872   1% /dev
devshm                  257968         0    257968   0% /dev/shm
192.228.6.64:/home/wvh
                     220010496 151044928   8965568  95% /mnt/nfs
```

The mounted directory is now available so that you can access the files that it contains and create new files there. Of course, because NFS filesystems are Linux (or UNIX) directories, you must have sufficient permissions to read and write files there. In most NFS environments, user identities are synchronized across systems so that you appear to be the same user on both the remote and local system. If this is not the case, you can always use the `sudo` command (if necessary) to copy files or directories to the NFS filesystem, as in the following example:

```
$ sudo cp resume.xml /mnt/nfs
```

Synchronizing user identities across multiple systems and understanding how Linux systems recognize users in general is discussed in Bonus Chapter 4, "Managing Users, Groups, Authentication, and Advanced Permissions."

Once you're done accessing the files in an NFS filesystem, you can unmount it using the standard Linux `umount` command, as in the following example:

```
$ sudo umount /mnt/nfs
```

Note that this command isn't "unmount," but umount. Why type that extra "n" if you don't have to?



If you want to automatically mount an NFS filesystem on your system, you should create an entry for it (filesystem type nfs) in your `/etc/fstab file`, as described in the section of Bonus Chapter 14 entitled "Mounting NFS Clients."

Peer-to-Peer File Sharing on Ubuntu

Let's say that you have a tarball or ISO of the greatest open source software package ever, everyone who hears about it wants a copy, and you want to make it freely available to all users. The classic solution to this problem was for you to put it on your FTP or, more recently, Web server, and then to advertise its existence to the known universe. Regardless of how fast your servers are, how many you have, and how fast your connection to the Internet is, it's easy to spot the problem with this approach: Eventually too many people show up to download the file, consuming all of the resources on your site until new people who've just heard about your whizzy software package can't get to it anymore, the delivery of your files to the people who are already in the process of downloading them slows to a crawl, and your servers fall over from the load. It's the downloading equivalent of being mentioned on SlashDot (`www.slashdot.org`).

The key to this problem is that there is only one place to get the magic files and everyone who wants them has to get them there. The entire load for grabbing these magic files is dumped on that single site, which turns it from a one-stop-shopping experience into a huge bottleneck. Even if you have an arsenal of replicated servers serving up your files and you have the world's best load-balancing software managing the traffic, your resources are finite, while the load may be effectively infinite.

Peer-to-peer (P2P) file sharing is an impressively cool solution to this problem. As mentioned in the introduction to this chapter, peer-to-peer (P2P) file sharing differs from traditional point-to-point file transfers in that P2P file transfers are essentially a community effort, a cooperative distribution network in which systems (peers) share files with each other without the user having to know which specific systems are involved in the transfer. Obviously, somebody has to know what hosts are sharing what files, or else no host would know where to get anything from, but there are some cool solutions to this problem. The legal Achilles heel of the pioneering P2P mechanism, Napster, was that it used a central database to track which peers were offering which files. Some legal idiots decided that this meant that Napster was effectively condoning all of the P2P traffic on the network, some of which involved sharing audio files that users might otherwise have had to pay for. They, therefore, forced Napster to shut down its central server, which effectively decapitated the original Napster network. Alas, poor Napster v1, I knew thee well.

Fortunately, performing a computer science autopsy of this "legal" fiasco has helped the second wave of P2P solutions come up with some great workarounds. These primarily revolve around peers dynamically monitoring and querying other peers or a network in its entirety to find out who has the files that you're looking for. The key to the success of the next wave of P2P solutions has been, and will continue to be, the effective anonymity of which peer is sharing what with whom. Without directly querying each and every peer, you can never get a complete picture of a modern, anonymous P2P network. Which is a good thing.

Today, there are many popular P2P protocols, file-sharing networks, protocols, and clients. Some of the most popular are the following:

- **BitTorrent:** www.bittorrent.com
- **eDonkey2000:** www.edonkey2000.com
- **FastTrack (essentially Kazaa):** www.kazaa.com
- **Freenet:** http://freenet.sourceforge.net
- **Gnutella:** www.gnutella.com

The next two sections explain how to be part of the free speech revolution by installing and using P2P clients such as `gtk-gnutella`, a popular Linux client for the Gnutella file-sharing network, `gnome-btdownload`, a simple but complete Linux client for BitTorrent downloads. As in the rest of this chapter, these sections focus on using client software to retrieve files, not on how to explicitly publish files from your system—you're sharing them while this is happening, of course, but the focus here is on retrieving them.

Installing and Using gtk-gnutella

The Gnutella network (not to be confused with Nutella, the yummy European hazelnut and chocolate spread) is a widespread, immensely popular P2P network. Many people do not realize that popular P2P applications such as LimeWire and BearShare are actually just Gnutella clients. One of the coolest things about Gnutella clients is their search feature, which I'll explain later and which makes it easy to query the entire network for shared files that match your search terms. As explained in the introduction to this section, the search feature does not consult a single centralized database because that would make it too easy to terminate Gnutella traffic. Instead, it uses a very cool friend-of-a-friend mechanism. Your query initially goes to a single Gnutella node, which passes on your original query to the other nodes it knows about, each of which passes on your original query to the other nodes it knows about, until much of the Gnutella network is buzzing with your search request.

Installing gtk-gnutella

Because not everyone wants or needs to do P2P file sharing, the `gtk-gnutella` client isn't installed by default on an Ubuntu Linux system. To install `gtk-gnutella`, use any of the package management tools discussed in Chapter 23 to install the `gtk-gnutella` package and any of its dependencies. Kubuntu users should note that installing `gtk-gnutella` will install a fair number of GNOME-related packages on your system if they're not already present.

Using gtk-gnutella

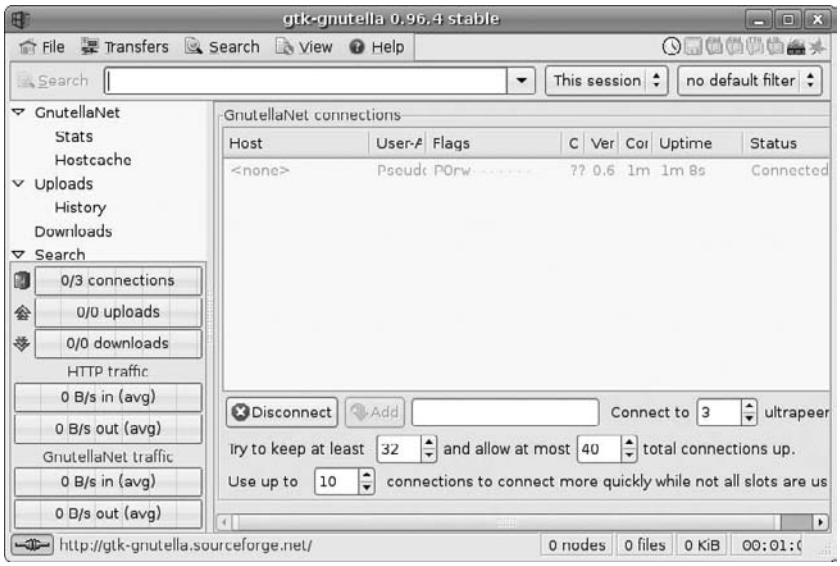
Installing gtk-gnutella as described in the previous section creates a menu item on the Applications ➪ Internet menu. To start gtk-gnutella, click the Gtk-Gnutella item on this menu, which starts the application and displays a window like the one shown in Figure B2-23.

The first time that you start gtk-gnutella, it will create the gtk-gnutella-downloads directory in your home directory, which is where it will store the files that it downloads. This directory contains three subdirectories: `complete`, which is where successfully downloaded files are stored when a download completes; `incomplete`, which is where in-progress downloads are stored; and `corrupt`, which is where downloaded files that do not match Gnutella's official checksum for the files are stored. Files in the `corrupt` directory are probably bad, but just in case. . .

NOTE **When you start `gtk-gnutella`, you may also see a dialog entitled "Ancient version detected!" that notes that the version of `gtk-gnutella` that is provided in the Ubuntu repositories isn't the latest version, and points you to `gtk-gnutella`'s SourceForge site for the latest version. Though this is true, I wouldn't bother to get the latest version from SourceForge unless you're actually having a problem with the version that is provided in the Ubuntu repositories. Not only will you have to build it yourself, but you won't be able to automatically upgrade it as you would if you used the standard version from the Ubuntu repositories and a new version was released there.**

FIGURE B2-23

The gtk-gnutella startup window



Aside from the standard menu bar, the most interesting parts of the gtk-gnutella screen shown in Figure B2-23 are the following:

- **graphical summary area (lower left):** An area that displays status information about the number of connected clients, the number of uploads and downloads in progress, the local HTTP (Web) traffic that the application sees, and the amount of Gnutella traffic it sees passing by on the network at the moment.
- **details window (upper left):** A window containing collapsible/expandable entries that enable you to get detailed information about any of the items that it lists: the state of the network, active and completed uploads, active and queued downloads, and the state of any searches that you have initiated.
- **search bar (upper left, just below the menu bar):** Enables you to search for files whose names match a specific sequence of characters.
- **connection status window (middle right):** Shows the status of any active connections, the results of a search when performing a search, and the status of any active downloads when you view active or queued downloads. The type of information displayed in this window depends on the item selected in the control window.
- **control window (bottom right):** Enables you to control the number of connections that are supported to and from your machine.

The interesting parts of gtk-gnutella are actually searching for something, seeing what matches you find, and then downloading the result. So let's first start a search. To demonstrate that there is plenty of legal and useful content out there, search for "Linux ISO." To do this, put the cursor in the search field of the

search bar, enter **Linux ISO**, and press return. The search begins, with any matches for the search string displaying in the connection status (now a search results) window. Once search results are displayed, you can see summary information about any of them by positioning the cursor over its name.



If you do not see results after a while, you may actually need to get a newer version of `gtk-gnutella` because the Gnutella network often bans sufficiently old versions from accessing the network. You can find the latest versions of `gtk-gnutella` for i386 and 64-bit systems at `http://archive.ubuntu.com/ubuntu/pool/universe/g/gtk-gnutella`, and manually install these using the `dpkg -i PACKAGENAME` command.

Once you've spotted a file that you want to download, you can select it for downloading by right-clicking on its name and selecting Download selected files from the pop-up menu.

Once you've started a download, you can view its status at any time by selecting the Downloads ➪ Active entry in the details window. If the download is not listed there, it is queued, and you can select the Downloads ➪ Queue entry in the details windows to see what it is waiting for.

Many systems in the Gnutella network offer many files for downloading and will therefore match search requests, but may already be delivering as many downloads as they have been configured to deliver by their operators. You may have to wait a while for a download to move from queued to active status. Frankly, if I find several matches for a specific file that I'm looking for, I often begin downloading more than one of them to ensure that I can get the file as quickly as possible. Once one of them actually begins to download, you can terminate the others so that you don't waste Gnutella bandwidth.



If you have problems getting `gtk-gnutella` to successfully locate or download files, make sure that any firewall that you are running does not block the ports that are used by the version of Gnutella. If you are seeing performance problems and are behind a configurable gateway or router, you may want to forward those ports from your external Internet interface directly to the host on which you are running your BitTorrent client. See the gtk-gnutella FAQ at `http://gtk-gnutella.sourceforge.net/en/?page=faq` for more information. A general discussion of port forwarding on a variety of routers is available at `www.portforward.com/routers.htm`.

Using BitTorrent

BitTorrent, originally written by Bram Cohen, has been around for a few years and is an incredibly popular protocol and tool for P2P file sharing. Bram's original BitTorrent software is GPL and was written in Python. Frankly, I still primarily use the command-line version of BitTorrent (btdownloadheadless.py) because "it just works." Don't worry, I'm not going to make you follow my lead here—although you can easily install the command-line version of the Python BitTorrent software by installing the bittorrent package. Instead, easy-to-use graphical BitTorrent clients are preinstalled on both Ubuntu and Kubuntu systems, which are what I'll focus on in this section.

The special sauce in the BitTorrent P2P mechanism is the use of torrent files, which provide information about the files that are to be distributed as part of a *torrent* (the generic BitTorrent term for each collection of files that are shared as a group), and the use of a computer system that tracks the current activity related to a torrent (known to BitTorrent devotees as a *tracker*). Unlike Napster's old use of a central server, a BitTorrent tracker doesn't know anything about the contents of a torrent but simply knows what peers are currently involved in downloading, uploading, and sharing a particular torrent. The tracker initially choreographs communications between any peers interested in a specific torrent—once peers are communicating with each other, the tracker is no longer necessary. As you'd expect, when initiating a torrent, at least one peer must provide a complete set of the files in the torrent, which is known as *seeding* in BitTorrent vernacular.

One interesting optimization in BitTorrent is that peers retrieve and share the files in a torrent in a random sequence of chunks. This means that, while many peers may be in the process of downloading the collection of files that make up a torrent, it's normal to expect that no two of them are sharing the same portions of those files. This makes torrents quite durable—even if the tracker and the host that is seeding a torrent goes away, it's quite possible that all of the peers together have a complete set of the files in the torrent—they all just have to trade missing chunks with each other. This is known as a *distributed copy* in BitTorrent terms.

BitTorrent is a great piece of software, an impressive protocol, an excellent software distribution mechanism, and a good example of how P2P file sharing can help guarantee the free sharing of data of any kind over the Net. It is also free. However, if you find it useful, you may want to consider donating funds to help guarantee its continued existence and development. I have. You can contribute to the future of free speech and free data on the Net via PayPal at `http://bittorrent.com/donate.myt`.

TIP **If you have problems getting any BitTorrent client to work, make sure that any firewall that you are running does not block TCP ports 6881-6999, which are the ports that are used by the current version of BitTorrent. If you are seeing performance problems and are behind a configurable gateway or router, you may want to forward those ports from your external Internet interface directly to the host on which you are running your BitTorrent client. Michael Ingram has a great article on solving slow downloading problems and configuring specific routers and gateways at `www.slyck.com/news.php?story=493`. A general discussion of port forwarding on a variety of routers is available at `www.portforward.com/routers.htm`.**

Using Transmission on Ubuntu Systems

Because Transmission is installed as part of a basic Ubuntu distribution, you don't have to do anything special to access or configure it. To start Transmission, select the Applications ➪ Internet ➪ Transmission BitTorrent Client menu command. This will display the initial dialog shown in Figure B2-24.

FIGURE B2-24

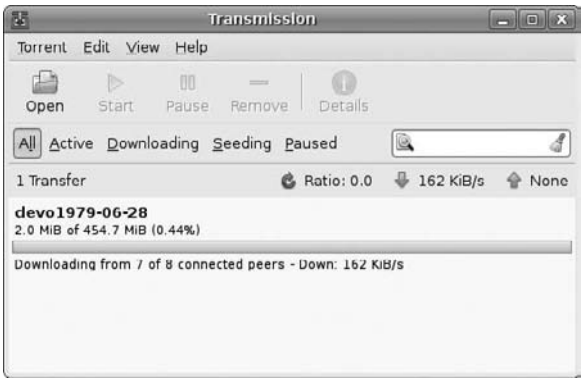
The initial Transmission dialog



Click the Open icon (or select the File ➪ Open menu command) and use the browse dialog to navigate to the directory that contains the torrent files that you have downloaded, select one, and click Open to begin downloading and sharing its contents. The Transmission dialog will show the status of your download, as shown in Figure B2-25 displays.

FIGURE B2-25

Downloading in progress in Transmission



The information under the name of the torrent but above the progress bar gives the amount of the torrent that you've currently downloaded, the complete size of the files in the torrent, and the percentage of the files in the torrent that you've currently downloaded. The number of peers from which you're downloading portions of the torrent displays below the progress bar.

Once you have downloaded the complete set of files that are found in a torrent, Transmission will begin to seed the torrent, which means that you are offering the files in the torrent to other people who are trying to download it. It's generally good form to seed a torrent until the amount of data that other users have downloaded from your seed matches the size of the torrent (as long as other people are trying to download the torrent—there's no need to wait forever). File sharing does indeed mean sharing and, as Mr. Rogers always said, "It's nice to share." Who would have thought that he was a BitTorrent fan?

TIP

If you're in the middle of downloading or seeding a torrent and you need all of your network bandwidth for some higher-priority task like transferring a file from work or watching something critical on YouTube, you can temporarily suspend BitTorrent activity by selecting each torrent and clicking Pause in the Transmission toolbar. When you're done with whatever was more important, you can resume downloading and/or seeding by selecting each paused torrent that you want to resume and clicking Start in the Transmission toolbar.

Transmission enables you to download and seed multiple torrents at the same time. By default, the Transmission dialog shows the status of all torrents, regardless of whether you're downloading or seeding them. To restrict the torrents shown in Transmission to those that are in a specific state, you can click any of the buttons above the status window. Clicking Active shows all non-paused torrents. Clicking Downloading shows all torrents that you are still downloading, while clicking Seeding shows all torrents whose downloads have completed but that you are providing to other people. As you might guess, clicking Paused only shows torrents that you have explicitly paused.

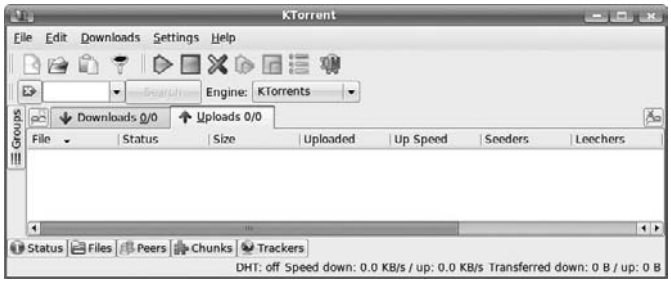
Happy file sharing! If you find the Transmission interface too minimal for your tastes, a variety of other graphical BitTorrent clients are available in the Ubuntu repositories. Some of the best known of these are `azureus`, a Java-based client that also includes a tracker and can therefore be used to seed torrents; `bittornado-gui`, a fast and powerful but relatively lightweight client; and `qtorrent`, a graphical client that uses the QT graphical interface.

Using KTorrent on Kubuntu Systems

Because KTorrent is installed as part of a basic Kubuntu distribution, you don't have to do anything special to access or configure it. To start KTorrent, select the K Menu ➪ Internet ➪ KTorrent - BitTorrent Client menu command. This will display the initial dialog shown in Figure B2-26.

FIGURE B2-26

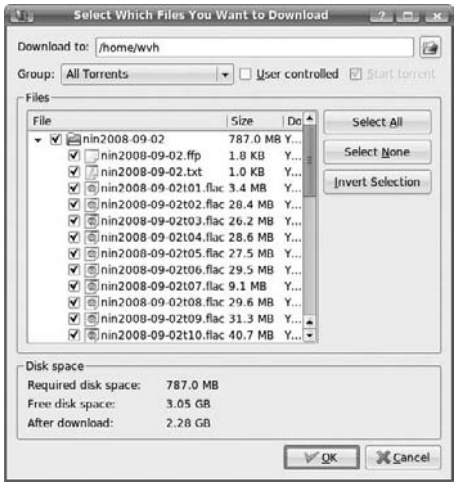
The initial KTorrent dialog



Click the Open icon (or select the File ➪ Open menu command) and use the browse dialog to navigate to the directory that contains the torrent files that you have downloaded, select one, and click Open to display a dialog that shows the contents of the torrent and the default location to which it will be downloaded, as shown in Figure B2-27.

FIGURE B2-27

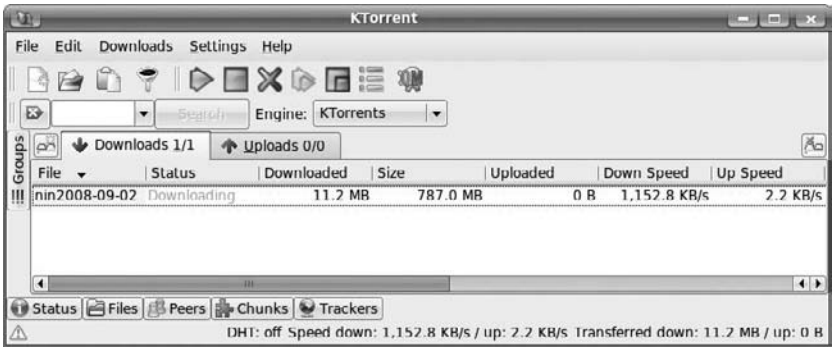
Information about a selected torrent



If you want to download the torrent to a specific location other than your home directory, you can use the Browse button to the right of the default download location at the top of this dialog to navigate to the directory where you want to store the directory of files that are associated with this torrent. If you only want to download specific files from the torrent, you can deselect the checkbox to the left of any of the files that you do not want to download. Once you have made any changes that you want to make, click OK to begin downloading the files associated with the torrent. You can then click the Download tab to see the status of the torrent(s) that you are downloading, as shown in Figure B2-28

FIGURE B2-28

Downloading in progress in KTorrent



The information about the torrent gives the amount of the torrent that you've currently downloaded, the complete size of the files in the torrent, your approximate download speed, the number of people that are currently seeding the torrent and the number of people that are currently downloading it (known as *leechers* in BitTorrent-speak), the percentage of the files in the torrent that you've currently downloaded, the ratio of data that you've downloaded to that which you've uploaded (if the latter is nonzero), and the elapsed time for which you've been downloading or seeding the torrent.

Once you have downloaded the complete set of files that are found in a torrent, KTorrent will begin to seed the torrent, which means that you are offering the files in the torrent to other people who are trying to download it. KTorrent enables you to download and seed multiple torrents at the same time, which you can see by clicking the Downloading or Uploading tabs, respectively. When you begin seeding a torrent, it is removed from the Downloading tab and displayed on the Uploading tab. It's generally good form to seed a torrent until the amount of data that other users have downloaded from your seed matches the size of the torrent (as long as other people are trying to download the torrent—there's no need to wait forever). File sharing does indeed mean sharing and, as Mr. Rogers always said, "It's nice to share." Who would have thought that he was a BitTorrent fan?

Finding Torrents

At this point, I hope that you're thinking, "Well, BitTorrent sounds like the best thing since sliced bread, but where do I get these torrents from?" Because BitTorrent doesn't provide the same sort of search features as P2P networks such as Gnutella, you have to have a torrent file to actually use it. You may have

noticed that the torrent I used as an example in the previous section is a live recording by the Grateful Dead, who are well known for allowing fans to freely share their live concert recordings. A quick search on the Net for "Grateful Dead" and "torrent" turns up zillions of hits, but that's hardly the most efficient search mechanism when looking for a live recording by your favorite band or some other freely available downloads.

Luckily, many sites provide search engines that are specifically designed for finding torrents. Some of my favorites are the following:

- **btjunkie:** `http://btjunkie.org`
- **isoHunt:** `http://isohunt.com`
- **Torrent Scan:** `www.torrentscan.com`

When getting started, you may also want to visit `http://netforbeginners.about.com/od/peersharing/a/torrent_search.htm`, which provides a user-submitted list of popular BitTorernt search engines.

Because BitTorrent is such an efficient P2P mechanism, many Linux distributions (including Ubuntu) provide torrents for their most recent distributions as well as supporting the standard, slow Web-based downloads. BitTorrent is a fast way of downloading the CD or DVD image (ISO file) for your favorite Linux distribution—which is Ubuntu or Kubuntu, I hope. You can find links to sites that provide torrents for the latest Ubuntu releases at `www.ubuntu.com/download`, and Kubuntu information at `www.kubuntu.org/getkubuntu`.

Getting More Information about BitTorrent

Because this is a general-purpose book about the Ubuntu and Kubuntu Linux distributions and not a BitTorrent bible, this section has only scratched the surface of available information about BitTorrent. I'm betting that once you've started to use BitTorrent, you'll have additional questions or will simply want to know more. Some great sources of dedicated information about BitTorrent are the following:

- **Bit Torrent Page:** `http://wiki.etree.org/index.php?page=BitTorrent`. A great collection of links to BitTorrent information on the Web from a popular filesharing site for legal, lossless audio recordings.
- **BitTorrent FAQ Guide Help:** `www.vladd44.com/torrent`. A useful guide to using BitTorrent and resolving common problems.
- **Brian Dessent's FAQ:** `www.dessent.net/btfaq`. Though somewhat dated in parts, this FAQ provides a good deal of useful information.

Summary

Business and academic environments typically provide file servers where centralized collections of corporate documents, project directories, and even supported software are stored. Many networked home computing environments do the same thing to provide a central location for family photos, music, and documents. This chapter began by explaining how to access remote file servers running on Microsoft

Windows, Mac OS X, and Linux systems using software such as FTP and how to mount shared folders and network directories using Samba and NFS. The last section of this chapter discussed peer-to-peer (p2p) file sharing on Ubuntu, the most popular and resource-friendly mechanism for sharing files that is in common use today.

Bonus Chapter 3 moves into territory that is the traditional bugaboo of Linux systems —how to compile your own software on a Ubuntu or Kubuntu system. Ubuntu and Kubuntu systems do not include a compilation environment in their default installations because not everyone needs them, but more importantly because the Ubuntu repositories provide such a great selection of software that you rarely have to do this unless you want to (or are a developer). The need to roll your own versions of applications on Linux systems is a traditional part of the FUD (Fear, Uncertainty, Doubt) that Linux detractors use to support their claims that Linux is hard to use and only for nerds. This is not true at all. Not only is building your own applications optional nowadays but, as you'll see in Chapter 23, package management software like `apt-get`, `aptitude`, `Adept`, and `Synaptic` make it easy to install build tools and a powerful IDE if you need them.

Bonus Chapter 3

Software Development on Ubuntu and Kubuntu

As pointed out on the Web site and throughout this book, Ubuntu and Kubuntu are distributions for everybody. And that includes developers! Just because Ubuntu and Kubuntu are easy to use doesn't mean that serious, hard-core developers can't get real work done on these systems—quite the contrary, actually. A good, recursive proof of this is the Ubuntu and Kubuntu distributions, which are built by Ubuntu and Kubuntu developers on Ubuntu and Kubuntu systems. In addition, the fact that Ubuntu and Kubuntu offer such a great user experience means that the time that you might ordinarily spend tweaking and fine-tuning your Linux desktop and installed software can be devoted to more productive pursuits, such as working on your personal code projects. The whole world may be waiting for your software!

Linux, in general, has a bad rep among many of the users of other platforms such as Microsoft Windows and Mac OS X, who often see Linux users as cave dwellers using VT100 terminals and living at the command line. Although there's a lot to be said for the command line, modern Linux distributions such as Ubuntu and Kubuntu strike just the right balance between providing graphical utilities for just about everything, with command-line access when and if you need it. The modern Linux model is strikingly similar to the user model provided by Mac OS X, and recent versions of Microsoft Windows include "new" tools that provide support for extensive use of the command line and sophisticated command-script support. Does any of that ring a bell?

This chapter explains how to install the packages that you'll need for most Linux software development projects, starting with the fundamentals of installing the world's best compiler (`gcc`, the C compiler that is provided as part of the Free Software Foundation's GNU Compiler Collection, GCC), related tools that you'll need to provide a complete software build environment on your Ubuntu or Kubuntu system, and then branches out to discuss other software packages that many people use when creating and maintaining software on Linux systems. These other packages include integrated development environments (IDEs) that put an easy-to-use face on modern software development, and source code control and management software so that you can easily track changes to your software.

IN THIS CHAPTER

Installing and using GCC

Using make to simplify recompilation

Building open source software

Installing and using the Eclipse IDE

Installing, configuring, and using Subversion

Installing and Using Development Software

Ubuntu and Kubuntu systems don't provide a software build environment out of the box because many people don't want to do software development. After all, these distributions provide a great user solution and a tremendous selection of software packages, so many people will never need to build their own software. Providing the software that users need, and providing it in an integrated, easy-to-install, and easy-to-manage way, is what Linux distributions are all about in the first place.

For me, the most important part of the term *software* is the word *soft*. Software is flexible. Different people need different software packages to reflect their interests. When you're writing your own program to do "the right thing" that no one else seems to understand the need for, or simply want to use a special-purpose package that you can't live without from some random Web site, you may need to compile software yourself.

The traditional Linux build environment is composed of several different software packages, providing compilers, a linker, an assembler, an archiver for building and working with libraries, a utility to simplify building complex software packages, and much more. As you'd expect, Ubuntu and Kubuntu make it easy to install all of these core components as one simple über-package. As explained in the next section, installing this single package pulls together all of the basic ingredients of a Linux build environment, which you can subsequently augment with any other favorite tools that you've come to know and love when building software on your Linux system.

Installing GCC Compilers, make, and Friends

The Ubuntu repositories that are shared by both the Ubuntu and Kubuntu distributions make it especially easy to install the core elements of a Linux build environment by providing a package called `build-essential` that pulls in the GCC C and C++ compilers, mandatory libraries, the `make` utility, and the current set of Linux kernel headers. You install this package on your system using the Synaptic Package Manager, Adept Manager, or the command-line `aptitude` or `apt-get` commands as described in Chapter 23, "Adding, Removing, and Updating Software."

Depending on the build requirements of the software or package that you are compiling, you may find that you will subsequently need to install other packages. For example, configuration-related utilities such as `autoconf` and `automake` are not included in the `build-essential` package. Similarly, the `build-essential` package only provides the GCC C and C++ compilers—if you want to compile applications written in any of the other languages supported by GCC, you will have to install those compilers and any associated libraries separately. These additional packages can always be installed, as needed, in the same way as the `build-essential` package.

GCC Compiler Overview

In uppercase letters, GCC is an acronym for the GNU Compiler Collection, formerly known as the GNU Compiler Suite. In lowercase letters, gcc is a specific compiler from the GNU Compiler Collection, the GNU C Compiler. GCC is a collection of compiler front ends for a common back-end compilation engine. The list of compilers provided by GCC includes C (`gcc`), C++ (`g++`), Objective C (`gcc` with a special library), Fortran (`gfortran`), and Java (`gcj`). GCC also has front ends for Pascal, Modula-3, and Ada 9x. The C compiler speaks several different dialects of C, including traditional and ANSI C. The C++ compiler is a true native C++ compiler—in other words, it does not first convert C++ code into an intermediate C representation before compiling it, as older C++ compilers did. GCC's C++ compiler, g++, creates native executable code directly from the C++ source code.

GCC is a GNU Project, which is directed by the Free Software Foundation (FSF). The FSF holds the copyright on the compilers and licenses the compilers under the terms of the GPL. Either individuals or the FSF hold the copyrights on other components, such as the runtime libraries and test suites, and these other components are licensed under a variety of free software licenses. The FSF also handles the legal concerns of the GCC project.

Retrieving Source Packages

Throughout this book, the information on installing software packages has always focused on retrieving and installing binary packages, which are the packages that contain executable programs, related precompiled libraries, and associated configuration data that may be either text or binary. Because this chapter discusses installing build tools, it seems reasonable to at least touch on how to retrieve the source code associated with any Ubuntu or Kubuntu binary package. In a nutshell, there are essentially two steps to this, both of which are most easily done using the command-line `apt-get` utility:

1. Retrieve the source code for the package that you want to build.
2. Satisfy any dependencies on other software, libraries, and so on that are required to build a specific package.

To retrieve the source code for any Ubuntu or Kubuntu package, use the `apt-get source package` command, where `package` is the name of the package that you want to build from source code. To install any other tools, libraries, header files, and so on that are required to actually build that package from its source, you use the `apt-get build-dep package` command, where `package` is the name of the package that you want to build.

See Chapter 23, "Adding, Removing, and Updating Software," for a more detailed discussion of satisfying build dependencies and retrieving package source code using the `apt-get` utility.

Using GCC's C Compiler

The `gcc` compiler supports a tremendous number of command-line options to control its behavior. It accepts both single-letter options, such as -o, and multiletter options, such as -ansi. Because it accepts both types of options, you cannot group multiple single-letter options together, as you may be used to doing with many GNU and UNIX/Linux programs such as ps, ls, and so on. For example, the multiletter option -pg is not the same as the two single-letter options -p -g. The -pg option creates extra code in the final binary that outputs profile information for the GNU code profiler, `gprof`. On the other hand, the -p -g options generate extra code in the resulting binary that produces profiling information for use by the prof code profiler (-p) and causes `gcc` to generate debugging information using the operating system's normal format (-g).

Despite its sensitivity to the grouping of multiple single-letter options, you can freely mix the order of options and compiler arguments on the `gcc` command line. The following two commands are identical in terms of what they tell `gcc` to do:

```
$ gcc -pg -fno-strength-reduce -g hello.c -o hello
$ gcc hello.c -o hello -g -fno-strength-reduce -pg
```

In most cases, the order of options and their arguments do not matter. In some situations, order does matter, especially if you use several options of the same kind. For example, the -I option specifies the directory or directories to search for include files. So, if you specify -I several times, `gcc` searches the listed directories in the order specified.

Compiling a single source file, hello.c, using `gcc` is easy—just invoke `gcc`, passing the name of the source file as the argument, as shown in the following example:

```
$ gcc hello.c
```

By default, the result is an executable file named a.out in the current directory, which you execute by typing ./a.out. To define the name of the output file that `gcc` produces, use the -o option, as in the following example:

```
$ gcc hello.c -o runme
```

If you are compiling multiple source files using `gcc`, you can specify them all on the `gcc` command line, as in the following example, which leaves the compiled and linked executable in the file named showdate:

```
$ gcc showdate.c helper.c -o showdate
```

If you want to compile these files incrementally and eventually link them into a binary, you can use the –c option to halt compilation after producing an object file. You can then use a single `gcc` command to link them all together and specify the name of the output file, as in the following example:

```
$ gcc -c showdate.c
$ gcc -c helper.c
$ gcc showdate.o helper.o -o showdate
```

It should be easy to see that a project consisting of more than a few source code files would quickly become exceedingly tedious to compile from the command line, especially after you start linking in external libraries, adding optimization commands, and providing other miscellaneous `gcc` options. The solution to this command-line tedium is the make utility, which is discussed in the next section.

Creating Simple Makefiles

Since just after the dawn of UNIX time, most UNIX applications have been built using a rule-driven application called make, which invokes the compilers and other utilities necessary to produce an executable from all the components of a program's source code. The make program also serves as the cornerstone for building C and C++ applications on Linux and all other UNIX-like systems. The make program does the following:

- Enables developers to identify relationships and dependencies between the source modules, include files, and libraries that are required for successful compilation
- Specifies the sequence in which things must be compiled to build an application successfully
- Avoids recompilation by limiting compilation to the portions of an application that are affected by any change to source code or statically linked libraries

The make program was originally written by Stu Feldman at AT&T for an early version of UNIX. As a part of AT&T UNIX, the source code for the make program was not freely available, and it has therefore largely been replaced by the GNU make program. GNU make provides all of the features of the original

make program and a host of others. Subsequent references to the make program throughout this chapter, therefore, refer to GNU make.

The rules used by any make program are stored in text files named Makefile, which contain dependency rules and the commands necessary to satisfy them. As you might expect, the Makefiles for complex applications can themselves be extremely complex as well as platform-specific, because they need to invoke platform-specific libraries and utilities that may be located in different directories or have different names on different systems. The simplest possible Makefile is the one for a simple application that is compiled from a single source code file, such as the legendary C language "Hello, World!" application. The Makefile for this application looks like this:

```
hello: hello.c
gcc hello.c -o hello
```

The first line identifies the binary target (hello) that this Makefile builds, and the fact that this binary must be regenerated whenever the source file hello.c is modified. The second line (which begins with a tab) specifies the command necessary to produce that binary target. To invoke this Makefile, you simply execute the make command in the directory containing the Makefile. The make program automatically looks for a file named Makefile and follows any build rules that it contains.

TIP **The fact that all lines containing a build rule must begin with a tab is one of the most common problems with handcrafted Makefiles.**

Although this example shows the true basics of a Makefile, it is also a trivial example that doesn't show off any of the real benefits that Makefiles can provide in terms of simplifying and automating compilation and recompilation. Let's consider another example. In the previous section, one of the examples I used compiled a single binary from two object modules, using the following command lines:

```
$ gcc -c showdate.c
$ gcc -c helper.c
$ gcc showdate.o helper.o -o showdate
```

A Makefile for this application would look something like the following:

```
showdate: showdate.o helper.o
  gcc showdate.o helper.o -o showdate

showdate.o: showdate.c
  gcc -c showdate.c
helper.o: helper.c
  gcc -c helper.c
```

This example highlights how a Makefile can reduce the amount of recompilation necessary when you have only modified some of the files associated with a project. For example, this Makefile states that the showdate binary depends on the two object modules showdate.o and helper.o, and provides a rule for generating the binary by using `gcc` to link the two object files. It then continues to highlight that each of these object modules depends on its associated C language file, and provides a rule for using `gcc` to regenerate each of the object modules.

This sample Makefile can simplify recompilation if, for example, only one of the C files was modified. The Makefile would regenerate to object module for only the C file that has been modified, and then relink the existing object file with the new object file to regenerate the binary.

As your projects become larger, you will want to take advantage of other Makefile features, such as defining and using variables, using some truly cryptic symbols to simplify make rules, and so on. Some good sources for more advanced information about the make program and Makefiles are the following:

- `www.linuxdevcenter.com/pub/a/linux/2002/01/31/make_intro.html`: A nice article on using `make` from O'Reilly & Associates Linux Development Center
- `www.tldp.org/HOWTO/Software-Building-HOWTO-3.html`: The section on using `make` from the Linux Documentation project's guide to Linux software development
- `www.gnu.org/software/make`: The home page for GNU `make`, which also includes links to the official make documentation
- `www.gnu.org/software/make/manual`: The official online home of the documentation for GNU `make`, which is available in various formats

Building Open Source Software

In the primordial days of Linux, a huge stumbling block for its adoption was the fact that, while it was all free, using a Linux system was very much a roll-your-own endeavor. This led to the birth of Linux distributions, which precompiled a basic set of utilities and provided some way of installing them on a computer system—even if that was just a README file that explained what to do. Primal Linux distributions such as TAMU (Texas A&M University), SLS (Softlanding Linux Systems), and Yggdrasil are the forefathers of the Ubuntu or Kubuntu system that you are using today, where configuration, installation, maintenance, and access to additional software are all supported by whizzy graphics and basically "just work." Today, you rarely have to build anything yourself—but you never know.

As a long-time Linux user, I've come to expect or depend on certain pieces of Linux software. This is usually fine, but software ages like anything else, and occasionally is simply replaced by newer packages that do similar things, and usually more, by taking advantage of other newer packages, libraries, and so on. In this case, the older package is eventually dropped from inclusion in various distributions at some point. This was the case for me in Ubuntu 8.04, which dropped support for the graphical music player that I've used for years, the X MultiMedia System, known to its friends as XMMS. XMMS has been replaced by Audacious, a very similar but newer application (discussed in Chapter 19), and a new version of XMMS, cleverly named XMMS2, is under development. However, one of the best things about XMMS is its plug-in architecture, which enabled one to add support for many different audio file formats and associated codecs. Since I'm also a long-time music collector, I have a fair number of digital recordings in formats that aren't supported by Audacious or XMMS2, and it's nice to use a single audio player that just works with all of my online music in its native formats. So what's a girl to do? The obvious solution is to build XMMS and the plug-ins that I need myself, which provides an excellent example for this section.

NOTE **The XMMS package I'm using as an example in this section uses standard X window system and GTK 1 graphics (one of the reasons that it was dropped from Ubuntu 8.04). You can build this package in exactly the same way on a Kubuntu system but, as with any GTK application on a KDE-based system, you'll be installing lots of "alien" packages that are not normally found in a KDE environment.**

As with all modern software, building open source software packages can be somewhat complex. The next few sections show what can be involved in building an open source package from source code, provide examples of the types of problems that you may encounter when trying to build open source software

yourself, and explain how to recognize and resolve these problems. Along the way, these sections also highlight some of the different ways in which you can build other packages that are associated with the one that you're really after.

NOTE **The following sections on building open source software from the command-line assume that you have already installed the build-essential package that was discussed earlier in this chapter.**

XMMS Build Overview

If you're reading this section because you simply want to build XMMS with support for MP3, OGG, SHN, and FLAC files, rather than learning how to identify and resolve problems when building an open source application, the following is a simple list of the steps that you should follow, with no rationale for any of them:

```
$ sudo aptitude install libglib1.2-dev libgtk1.2-dev libesd0-dev \
                        libmikmod2-dev libogg-dev libvorbis-dev \
                        libasound2-dev libgl1-mesa-dev
$ cd ; mkdir src ; cd src
$ wget http://xmms.org/files/1.2.x/xmms-1.2.11.tar.gz
$ tar zxvf xmms-1.2.11.tar.gz
$ cd xmms-1.2.11
$ ./configure
$ sudo make install
$ cd ..
$ wget http://etree.org/shnutils/dist/src/xmms-shn/xmms-shn-2.4.1.tar.gz
$ tar zxvf xmms-shn-2.4.1.tar.gz
$ cd xmms-shn-2.4.1
$ ./configure
$ sudo make install
$ cd ..
$ sudo apt-get build-dep flac
$ apt-get source flac
$ cd flac-1.2.1
$ ./configure --enable-local-xmms-plugin
$ make
$ mkdir -p ~/.xmms/Plugins
$ cp src/plugin_xmms/.libs/libxmms-flac.so ~/.xmms/Plugins
```

If you simply follow these steps to build XMMS and are then curious about what each step actually did, you can always consult the appropriate section of this chapter for detailed information.

Building a Standalone Package

The Ubuntu repositories are great and contain most of the software that most people will ever want. However, if you're looking for a specific application and it isn't in the Ubuntu repositories, the first thing

to do is to use a search engine such as Clusty or Google to try to locate someone who has already built the application that you're looking for and offers it as a DEB package that you can install on your Linux system. Some trust issues are involved in this scenario—you have to hope that the person who built the package did it correctly, and also that they didn't hack the software in some evil way. The former is usually the case, and the latter has never happened to me, but you never know. So, it's far safer to simply build and install the package yourself.

NOTE **Building an open source application that is otherwise unavailable or unsupported also gives you the opportunity to be a good Ubuntu/Kubuntu citizen by offering it to others if you build it as a redistributable Debian (DEB) package and provide links to it on your Web site. While Ubuntu has some great tools for building DEB packages, many inherited from the Debian distribution, explaining how to do that is a topic with a fairly small audience and is, therefore, outside the scope of this chapter.**

Like many long-time UNIX/Linux users, I tend to build software in a subdirectory of my home directory called `src`. Regardless of whether you share my convention, it's a good idea to build software in subdirectories of a central location, so that you know where to look if you need to rebuild or share some package that you've written or are simply maintaining for yourself. A `~/src` directory isn't created by default as part of a user account on Ubuntu or Kubuntu systems, so for the purposes of this tutorial, you can create this directory and make it your working directory by using the following commands from an xterm, GNOME console, or KDE Konsole window:

```
$ mkdir ~/src ; cd ~/src
```

As they say in my favorite computer-related movie, the key to building software is to "Use the source, Luke!". Your favorite search engine should make it easy to locate the search with a query like "`package-name source download.`" In my case, the Clusty search engine identified that I could find the source code for XMMS at `www.xmms.org`, from the URL `http://xmms.org/files/1.2.x/xmms-1.2.11.tar.gz`. I would normally retrieve this using a browser, but for the purposes of this example, you could retrieve it by using the command-line `wget` command, as in the following example:

```
$ wget http://xmms.org/files/1.2.x/xmms-1.2.11.tar.gz
```

Once you retrieve the source package, you need to extract its contents, since the filename indicates that this is a gzipped tar file, you could extract its contents with a command like the following:

```
$ tar zxvf xmms-1.2.11.tar.gz
xmms-1.2.11/
xmms-1.2.11/intl/
xmms-1.2.11/intl/Makefile.in
[ much output deleted]
```

As you can see from this example, extracting a source tarball usually creates a directory that contains all the source code below it. You then make that directory your working directory, and the fun is almost ready to begin:

```
$ cd xmms-1.2.11
```

Many simple open source software packages with a limited number of dependencies on other software packages provide a Makefile (discussed in the previous section of this chapter) that you can use to build the software. More complex software packages, which probably includes almost anything with a

graphical component), use an applications called `autoconf` and `automake` to automatically generate a configuration script that serves a few purposes:

- Locate the software (include files, libraries, etc.) on which the package depends and codify those locations for the build process
- Identify any missing dependencies so that you have some idea of what's missing and can install it
- Create one or more Makefiles that you can subsequently use to build the software

You usually don't have to create the configure script yourself—that is usually done by the developers themselves. You simply have to run the configure script, watch for problems, and resolve them. Two basic classes of problems can be identified by a configure script:

- An unsatisfied dependency without which you cannot build the software at all. These should cause the configure script to exit without creating a Makefile.
- An unsatisfied dependency that simply disables some capability of the software. In other words, you can build the package without satisfying the dependency, but the software won't have all of the capabilities that it supports. These do not cause the configure script to exit and can easily get lost in the verbose output of the configure script unless you watch for them closely. Luckily, as shown later in this section, configure scripts for applications that can be built with different sets of capabilities usually summarize the capabilities with which the software will be built at the end of a run of the configure script. You can then look for the unsatisfied dependency, resolve it, and rerun the configure script to generate an updated Makefile.

Running the configure script for XMMS quickly identifies a missing dependency:

```
$ ./configure
checking build system type... i686-pc-linux-gnu
checking host system type... i686-pc-linux-gnu
checking for a BSD-compatible install... /usr/bin/install -c
[much output deleted]
checking for GLIB - version >= 1.2.2... no
*** The glib-config script installed by GLIB could not be found
*** If GLIB was installed in PREFIX, make sure PREFIX/bin is in
*** your path, or set the GLIB_CONFIG environment variable to the
*** full path to glib-config.
configure: error: *** GLIB >= 1.2.2 not installed...
```

Though the output of the configure script suggests that the version of the package it's looking for might be on your system but in some weird location, that's usually not the case. The final statement is usually true, so the easiest thing to do at this point is to figure out what package contains the missing software and install it. You can do this using `dpkg` or `aptitude`, as described in Chapter 23. The following example shows the results of a search from `aptitude`:

```
$ aptitude search glib1
i A libavahi-glib1          - Avahi glib integration library
p   libglib1.2-dbg          - The GLib library of C routines (debug)
p   libglib1.2-dev          - The GLib library of C routines (dev)
i A libglib1.2ldbl          - The GLib library of C routines
p   libxmmsclient++-glib1   - XMMS2 - glib client library for c++
p   libxmmsclient-glib1     - XMMS2 - glib client library
```

NOTE **Ironically, this search shows that some libraries for XMMS2 are available, but those aren't what we're interested in, since we're trying to build the original XMMS.**

The "p" in the first column for the `libglib1.2-dev` entry means that this package can't be found on our system. The `-dev` suffix indicates that this package contains items required for software development, typically both libraries and header files. Let's install this package:

```
$ sudo aptitude install libglib1.2-dev
Reading package lists... Done
Building dependency tree
Reading state information... Done
Reading extended state information
Initializing package states... Done
Building tag database... Done
The following NEW packages will be automatically installed:
  pkg-config
The following NEW packages will be installed:
  libglib1.2-dev pkg-config
0 packages upgraded, 2 newly installed...
Need to get 229kB of archives. After unpacking 729kB will be used.
Do you want to continue? [Y/n/?]
```

This output shows that installing the `libglib1.2-dev` package will also install the `pkg-config package`, which it depends upon. You can type `y` or simply press Return to proceed with the installation (the fact that the `Y` is capitalized in the prompt means that it is the default response if you simply press the return/enter key on your keyboard.)

After a lot more output, the package is installed, and you can try running the configure script again. If you're lucky, this will complete successfully, but in this particular example, you'll have to install another mandatory package that causes the configure script to terminate without generating a Makefile. This is the `libgtk1.2-dev` package, which you install in the same way as the `libglib1.2-dev` package. This package pulls in a large number of X Window system development packages, libraries, and so on. (No output shown.)

Once you've installed both the `libglib1.2-dev` and `libgtk1.2-dev` packages, rerunning the configure script will actually complete and generate a Makefile. However, looking at the summary displayed at the end of running the configure script suggests that we're not actually done:

```
$ ./configure
[ much output deleted ]
config.status: executing depfiles commands

Configuration:
  Install path:              /usr/local
  Build OSS plugin:          yes
  Build esd plugin:          no
  Build Solaris plugin:      no
  Build BSD Sun plugin:      no
  Build ALSA plugin:         no
  Build mikmod plugin:       no
  Build Ogg Vorbis plugin:   no
  Build OpenGL plugins:      no
```

```
  Pthread flag:               -lpthread
  Use one plugin dir:         no
  Allow user plugin dir:      yes
```

This output shows that none of the XMMS plug-ins that support both music formats other than MP3 and sound services other than OSS (Open Sound System) would be built if we built the software as it stands. Since this isn't the year 2003, this is clearly insufficient. How do we add support for those things? The key lies in examining the output of the configure script more closely. Scrolling back in the terminal window shows messages like the following:

```
checking for esd-config... no
checking for ESD - version >= 0.2.8... no
*** The esd-config script installed by ESD could not be found
*** If ESD was installed in PREFIX, make sure PREFIX/bin is in
*** your path, or set the ESD_CONFIG environment variable to the
*** full path to esd-config.
checking for libmikmod-config... no
checking for libmikmod - version >= 3.1.5... no
*** The libmikmod-config script could not be found
[some output removed]
checking for Ogg... no
*** Could not run Ogg test program, checking why...
*** The test program failed to compile or link...
[some output removed]
checking for Vorbis... no
*** Could not run Vorbis test program, checking why...
*** The test program failed to compile or link...
```

TIP **You can rerun the configure script and redirect its output into a file if you can't scroll back far enough in your terminal window to actually see all of its messages.**

These messages imply that the development packages for the ESD (old Linux Enlightened Sound Daemon), the MikMod library (for working with older sound files produced by Soundtracker, Protracker, and similar tools), OGG (an open source container format used in various open source audio, speech, and video codecs), and Vorbis (the open source audio codec most commonly used in OGG audio containers).

ESD is still used on some systems, and we need the OGG and Vorbis plug-ins, so you would use similar `aptitude` search commands as the previous ones to locate the development packages that you need, and similar `aptitude` install commands to install them. The installation command would be the following:

```
$ sudo aptitude install libesd0-dev libmikmod2-dev \
                        libogg-dev libvorbis-dev
```

After installing these packages, and rerunning the configure script, things are much better, but are not quite right:

```
$ ./configure
[ much output deleted ]
config.status: executing depfiles commands

Configuration:
  Install path:               /usr/local
  Build OSS plugin:           yes
```

```
  Build esd plugin:           yes
  Build Solaris plugin:       no
  Build BSD Sun plugin:       no
  Build ALSA plugin:          no
  Build mikmod plugin:        yes
  Build Ogg Vorbis plugin:    yes
  Build OpenGL plugins:       no
  Pthread flag:               -lpthread
  Use one plugin dir:         no
  Allow user plugin dir:      yes
```

Since most audio applications on modern Linux systems use the ALSA sound system at some level, this is clearly not sufficient. Identifying what's still missing requires closer inspection of the output of the configuration script. Searching for ALSA shows the following:

```
checking for ALSA CFLAGS...
checking for ALSA LDFLAGS...  -lasound -lm -ldl -lpthread
checking for libasound headers version >= 0.9.0... not present.
```

That seems important. Using `aptitude` to search for `'asound'` reveals the `libasound2-dev` package. After installing this, things are even closer:

```
$ ./configure
[ much output deleted ]
config.status: executing depfiles commands

Configuration:
  Install path:               /usr/local
  Build OSS plugin:           yes
  Build esd plugin:           yes
  Build Solaris plugin:       no
  Build BSD Sun plugin:       no
  Build ALSA plugin:          yes
  Build mikmod plugin:        yes
  Build Ogg Vorbis plugin:    yes
  Build OpenGL plugins:       no
  Pthread flag:               -lpthread
  Use one plugin dir:         no
  Allow user plugin dir:      yes
```

We're still missing the OpenGL plug-in, which is necessary if you've become accustomed to displaying visual effects while playing music in XMMS. We're also quite close to building all default XMMS plug-ins, so let's see if we can find out what's missing. This one is trickier to find in the output, but searching for GL eventually shows the following messages:

```
checking for glBegin in -lGL... no
checking for glBegin in -lMesaGL... no
checking GL/gl.h usability... no
checking GL/gl.h presence... no
checking for GL/gl.h... no
checking GL/glx.h usability... no
checking GL/glx.h presence... no
checking for GL/glx.h... no
```

Using `aptitude` to search for `'gl'` doesn't show any dev packages, but searching for mesa, shown in the second message, suggests that `libgl1-mesa-dev` contains both `'gl'` and `'mesa'`. After installing this package, the output of the configure script shows that all of the default Linux plug-ins will be built:

```
$ ./configure
[ much output deleted ]
config.status: executing depfiles commands

Configuration:
  Install path:              /usr/local
  Build OSS plugin:          yes
  Build esd plugin:          yes
  Build Solaris plugin:      no
  Build BSD Sun plugin:      no
  Build ALSA plugin:         yes
  Build mikmod plugin:       yes
  Build Ogg Vorbis plugin:   yes
  Build OpenGL plugins:      no
  Pthread flag:              -lpthread
  Use one plugin dir:        no
  Allow user plugin dir:     yes
```

Since we're building on Linux, we don't care about the Solaris and BSD Sun plug-ins, so we can simply build and install the software using simple `make` commands:

```
$ sudo make all install
```

This builds everything that is listed in the configure script and then installs it. Explicitly specifying all is actually extraneous, since the install make target depends on the all make target, but I wanted to be clear that nothing has actually been built yet.

Once this command completes, XMMS is built and installed into `/usr/local/bin`, the default location for `autoconf/automake`-based software packages. You can change this location by providing options to the configure script—type `./configure --help` for a complete list of these. I generally prefer to install into the `/usr/local` hierarchy for two reasons:

- Finding software in `/usr/local` makes it clear that the software was compiled locally, so you probably have to manually install this software on other systems
- If your system already contains another version of the same software, installing into `/usr/local` prevents prebuilt software from being overwritten. Overwriting existing software can prevent updates from being applied to your system, which is a bad thing for other packages that may depend on shared libraries, include files, and so on.

The downside of installing into the `/usr/local` hierarchy is that applications that are installed into that location and which use shared libraries will not initially be able to find them. For example, trying to execute the XMMS binary in `/usr/local` bin will display an error message and exit, as in the following example:

```
$ /usr/local/bin/xmms
xmms: error while loading shared libraries: libxmms.so.1
      cannot open shared object file: No such file or directory
```

To resolve this problem, you must add `/usr/local/lib` to the list of directories in which your system looks for shared libraries. You can do this by using a text editor to edit the file `/etc/ld.so.conf` and adding the directory `/usr/local/lib` to the end of this file. After doing this, you can execute the `ldconfig` command (via sudo) to have your system add the contents of this directory (and any subdirectories that it contains) to the list of locations where it looks for shared libraries. You can then run `/usr/local/bin/xmms` with no problems.

Building software that uses shared libraries in nonstandard locations will display messages like the following during the installation process (that is, during the make install process):

```
----------------------------------------------------------------------
Libraries have been installed in:
   /usr/local/lib

If you ever happen to want to link against installed libraries
in a given directory, LIBDIR, you must either use libtool, and
specify the full pathname of the library, or use the `-LLIBDIR'
flag during linking and do at least one of the following:
   - add LIBDIR to the `LD_LIBRARY_PATH' environment variable
     during execution
   - add LIBDIR to the `LD_RUN_PATH' environment variable
     during linking
   - use the `-Wl,--rpath -Wl,LIBDIR' linker flag
   - have your system administrator add LIBDIR to `/etc/ld.so.conf'

See any operating system documentation about shared libraries for
more information, such as the ld(1) and ld.so(8) manual pages.
----------------------------------------------------------------------
Messages such as these identify the possibility that locating the
shared libraries for the application may be a potential problem, but
aren't exactly intuitive.
```

NOTE **The `/etc/ld.so.conf` file is automatically scanned whenever you boot your system. Once you've added an entry to this file, your system will always look there—you do not have to rerun the `ldconfig` command.**

Building Add-On Software

As mentioned earlier and shown in the previous section, XMMS provides a plug-in architecture that makes it easy to add libraries that provide support for other audio file formats. One that I use is Shorten (SHN), which is a free lossless audio codec that has been popular for years. The source code for this codec is freely available, and no package that supports this audio format is available on Ubuntu or Kubuntu systems by default.

The fact that no SHN-related software is provided in the Ubuntu repositories makes it very easy to build a SHN plug-in for XMMS, because there are no possible conflicts between installed versions of SHN-related software and what we're building. The next section discusses how to build and integrate software that is provided with Ubuntu with an unsupported package such as XMMS.

Because there are no possible conflicts, the process of building and installing the SHN codec for XMMS is the following:

```
$ wget \
```

```
  http://etree.org/shnutils/dist/src/xmms-shn/xmms-shn-2.4.1.tar.gz
[output deleted]
14:44:09 (251.02 KB/s) - `xmms-shn-2.4.1.tar.gz' saved [307397/307397]

$ tar zxvf xmms-shn-2.4.1.tar.gz
xmms-shn-2.4.1/
[output deleted]

$ cd xmms-shn-2.4.1

$ ./configure
checking build system type... i686-pc-linux-gnu
[output deleted]
configure: ready to compile.

xmms-shn is now configured with the following options:

version:      2.4.1
install dir:  /usr/local/lib/xmms/Input

$ sudo make install
[sudo] password for wvh:
Making install in utils
[output deleted]
make[1]: Leaving directory `/home/wvh/src/xmms-shn-2.4.1'
```

As you can see from this example, all of the requirements for this plug-in have already been satisfied by the requirements for XMMS itself, and this package can be configured, compiled, and installed with very little effort.

Things get a bit more interesting when a package that is provided in the Ubuntu repositories can also provide plug-ins or other software for use with packages that are not. This is discussed in the next section.

Building Supported Software in Different Ways

One of the other lossless audio codecs that I depend on is the Free Lossless Audio Codec (FLAC). FLAC support is provided in the Ubuntu repositories because it is used with packages such as GStreamer and others that were discussed in Chapter 19. However, the FLAC package that is provided in the Ubuntu repositories does not build the plug-in for XMMS, since XMMS itself is no longer supported. Therefore, we'll need to build the FLAC package in a different way than it is ordinarily built and install it so that it doesn't conflict with the binaries, libraries, and other files that are provided with the official FLAC package. This ought to be quite easy for several reasons:

- The Ubuntu FLAC package is installed under the `/usr` directory hierarchy, while the version of XMMS that we've built is installed under `/usr/local`.
- `/usr/local` is the default location for installing most open source software. This directory is overridden by `/usr` during the Ubuntu build process, so if we don't override the defaults, our custom FLAC software will automatically use `/usr/local`.
- It is easy to use package management commands to retrieve the source code for packages that are provided in the Ubuntu repositories, as discussed in the sidebar earlier in this chapter entitled "Retrieving Source Packages." No Web search necessary!

- It is easy to satisfy any build dependencies for packages that are provided in the Ubuntu repositories. This was also discussed in the sidebar earlier in this chapter entitled "Retrieving Source Packages."

All of that said, building the FLAC plug-in for XMMS can be done by executing the following short sequence of commands. There's no real need to show the output, since no interaction is required beyond these commands:

```
$ sudo apt-get build-dep flac
$ apt-get source flac
$ cd flac-1.2.1
$ ./configure --enable-local-xmms-plugin
$ make
```

Note that, though the software was configured to be installed under `/usr/local` by the configure script, the previous commands just built it—they did not install it. Though we could easily have done this automatically without compromising the FLAC software that is provided in the Ubuntu repositories, it is unnecessary to install everything, because we simply want to install the FLAC plug-in for XMMS, not everything in the package. After all, Ubuntu already provides everything else in its package!

To install just the XMMS plug-in, you can create a personal plug-in directory and install the plug-in there, as is done in the following example:

```
$ mkdir -p ~/.xmms/Plugins
$ cp src/plugin_xmms/.libs/libxmms-flac.so ~/.xmms/Plugins
```

The plug-in will now be available the next time that you start XMMS.

Using an Integrated Development Environment

An *integrated development environment* (*IDE*) is a unified set of tools that typically provides an editing environment, compiler, build environment, and debugger that work together in a single framework. IDEs typically also provide infrastructure for creating and managing software projects composed of multiple source files, include files, and libraries. Most modern IDEs are graphical tools that work in desktop environments such as the X Window system, GNOME, KDE, Microsoft Windows, or Mac OS X. For this reason, many modern IDEs, therefore, also include tools for creating user interfaces and the other graphical elements of graphical desktop applications.

The key aspect of an IDE is the integration among the different tools that make up the development environment that it provides. Most IDEs give the impression of being a single tool, regardless of how many features and components they provide.

As I'll discuss in the next section, IDEs are traditionally oriented toward the requirements of developers working in a specific language or a set of related languages. Recently, general-purpose IDEs such as Eclipse have become more popular, because they provide a single framework that standardizes the developer experience regardless of the language in which a developer is working, and support a plug-in model that enables them to easily be extended.

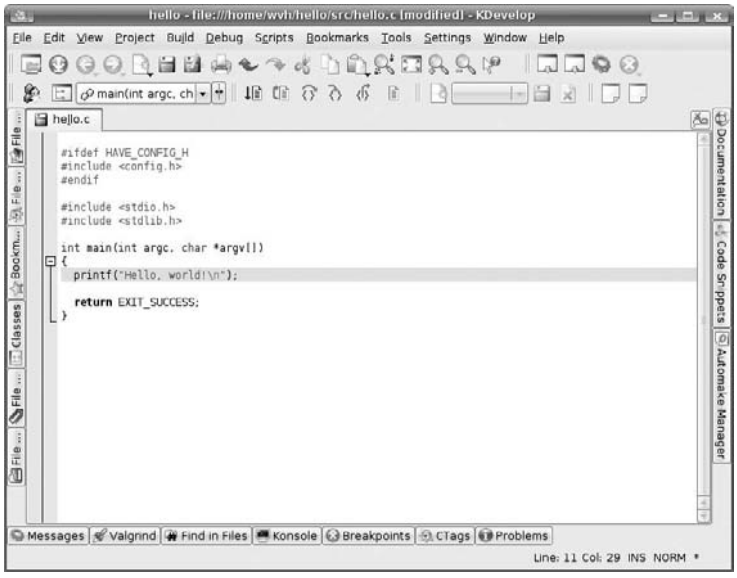
Popular IDEs for Linux

As you might expect from an operating system that is used for the development of applications in a huge number of languages, and one in which everyone has their own idea of how development should be done, Linux is knee-deep in IDEs. Ubuntu provides a large selection of IDEs for different languages in its repositories. Some of the more widely known IDEs provided in the Ubuntu repositories are the following:

- **Eclipse:** Eclipse is a multipurpose IDE written in Java. Eclipse was originally developed as an IDE framework for developing Java applications, but has been extended to support development in other languages. The fact that Eclipse is written in Java makes it very portable and usable in a variety of different graphical environments and operating systems. For more information about Eclipse, see the rest of the discussion of IDEs in this chapter or the Eclipse home page at `www.eclipse.org`.
- **KDevelop:** A graphical IDE for C and C++ applications that is based on the KDE desktop environment. KDevelop is an extremely popular and easy-to-use IDE that provides wizards, which simplify creating standard command-line Linux applications, as well as graphical applications that use the Qt, GNOME, and KDE graphical environments. For more information about KDevelop, see its home page at `www.kdevelop.org`. Figure B3-1 shows a version of the traditional C language "Hello, World!" application in KDevelop.
- **Motor:** The motor IDE is a text-based integrated development environment for C, C++, and Java applications that works well within any terminal application, such as an xterm or the GNOME Terminal. The motor IDE provides integrated support for various GCC compilers and the GNU debugger, `gdb`. For more detailed information about Motor, see its home page at `http://konst.org.ua/motor`.

FIGURE B3-1

A sample C application in KDevelop



- **NetBeans:** The NetBeans IDE is a complete IDE written in Java for developing applications in languages such as Java, C++, and Ruby. Because it is written in Java, it runs on many platforms including Microsoft Windows, Linux, Mac OS X, and Solaris, and is therefore a good choice for a cross-platform IDE alternative to Eclipse. For more detailed information on NetBeans, see its home page at `http://netbeans.org`.
- **Pida:** The pida IDE (Python Integrated Development Application) is a GTK-based IDE for developing text mode or graphical application written in the Python scripting language. The pida IDE is quite popular and supports easy integration of familiar development tools within its unified development framework. For more information about pida, see its home page at `http://pida.berlios.de`.
- **Sbcl:** The `sbcl` IDE (Steel Bank Common Lisp) is a graphical development environment for applications written in the Lisp programming language. The sbcl IDE was originally based on the Common Lisp development environment provided by Carnegie Mellon University's CMU Common Lisp, but is actively being maintained and enhanced. The `sbcl` IDE features provide a native Lisp code compiler and a sophisticated debugger. For more information about `sbcl`, see its home page at `http://sbcl.sourceforge.net`.
- **Free-java-sdk:** The Free Java SDK is an IDE created by combining and integrating best-of-breed open source tools to provide an alternative to the development environment provided by the standard Java Development Kit from Sun Microsystems, Inc. The free-java-sdk package was assembled from various open source projects by the developers of the Debian Linux distribution and has, therefore, been inherited by Ubuntu Linux.

This list is just a subset of the IDEs available in the Ubuntu repositories, but highlights the number of different IDEs available and the fact that many IDEs focus on specific programming languages. My favorite IDE is the Eclipse IDE, which was originally developed by IBM as an easily extended framework for Java development, but was released by IBM as open source, and has soared in popularity since then. Eclipse is now released and supported by the Eclipse Foundation (`www.eclipse.org`), a not-for-profit organization dedicated to promoting Eclipse and helping guide its development and acceptance throughout the software development industry. Eclipse provides a plug-in based framework that makes it easy to create, integrate, and support tools that provide specific functionality or target certain sets of developers. This plug-in model also gives developers the freedom to select and integrate plug-ins from different projects or commercial vendors and thereby assemble an IDE that does exactly what they need it to do. Eclipse is available for Linux, HP-UX, AIX, Solaris, QNX, Mac OS X, and Microsoft Windows systems.

As you might suspect from the previous paragraph, the rest of this part of the chapter focuses on installing and using the Eclipse IDE.

Installing Eclipse

To install Eclipse and related packages, use your favorite software installation tool to install the core eclipse package and any others that you might be interested in. Installing the core Eclipse package installs a huge number of associated/required packages, so I typically install the following packages:

- `eclipse:` (Mandatory) This is the core Eclipse package and provides everything you need to use the standard Eclipse IDE to develop Java applications, including plug-ins for Java and Python development, the GNU compiler for Java, the ant utility for building Java applications, all required libraries, and so on.

- `eclipse-cdt`: (Optional) Eclipse C Developer's Toolkit. This package is required if you want to build C or C++ applications in the Eclipse IDE. This package provides plug-ins, wizards, and Eclipse features to simplify developing and debugging C and C++ applications using Eclipse.

Once the installation completes, you're ready to begin developing applications using the most popular cross-platform IDE available today.

See Chapter 23, "Adding, Removing, and Updating Software," for detailed information on installing applications on Ubuntu and Kubuntu systems.

Using Eclipse for C Application Development

This section provides two examples of using Eclipse: one which shows how to create a simple application from scratch, and one that shows how to import the source code for an existing software package into Eclipse to take advantage of its compilation, debugging, and usability features. Entire books are available about using Eclipse, and these introductory sections are hardly equivalent to those, but they should help you get a feeling for the power, capabilities, and ease-of-use of the Eclipse IDE.

NOTE **The examples in the next two sections are both C programs, and therefore assume that you took my advice in the "Installing Eclipse" section and installed both the `eclipse` and `eclipse-cdt` packages, because the latter provides the support for C application development that these examples take advantage of. The examples in this section also assume that you followed the instructions in the "Installing GCC Compilers, make, and Friends" section, and installed the `build-essential` package, which provides the versions of `gcc` and `make` that Eclipse requires to build C applications.**

Starting Eclipse

Installing Eclipse on your Ubuntu system creates the Applications ➪ Programming ➪ Eclipse menu entry. Installing Eclipse on your Kubuntu system creates the K Menu ➪ Development ➪ Eclipse menu item. Select the appropriate one of these to start the Eclipse IDE.

The first time that you start Eclipse, Eclipse displays a dialog that prompts you for a location in which to create a workspace, and then creates the workspace and .eclipse directories in the specified location (usually your home directory).

TIP **Eclipse uses the workspace directory to store the projects that you create, and uses the .eclipse directory to store information about your configuration settings, the organization of the Eclipse UI the last time you used it, and so on. You can start Eclipse using a different workspace directory by using the -data command-line argument, followed by full pathname of the workspace directory that you want to use.**

After creating these directories, Eclipse then displays a Welcome screen, as shown in Figure B3-2.

The graphical objects on the Welcome screen are actually links to sources of information about working with Eclipse, such as tutorials, sample code, and so on. To close the Welcome screen, simply click the X in the Welcome tab. You can reopen it at any time by selecting Help ➪ Welcome.

After closing the Welcome screen, the standard Eclipse user interface displays, as shown in Figure B3-3.

At this point, the Eclipse environment is remarkably empty because you haven't created any projects. The next two sections describe how to fix that.

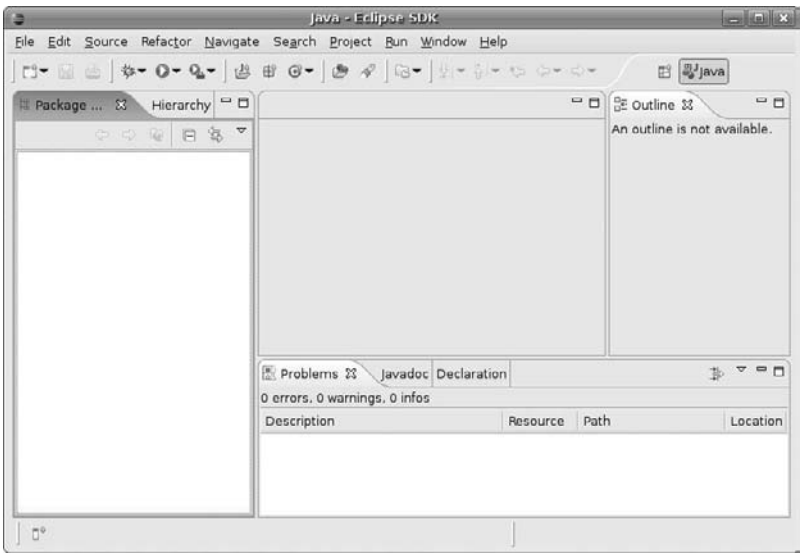
FIGURE B3-2

The welcome screen in Eclipse



FIGURE B3-3

The standard Eclipse user interface



Eclipse Terminology

The Eclipse framework was originally developed by IBM, which seems to have also developed its own language to refer to the elements of the Eclipse user interface and framework. This shouldn't be much of a surprise to anyone who has worked with IBM in the past. The following are a few terms that you'll encounter when using Eclipse—like most languages, they make perfect sense once you have been assimilated, but they can be confusing at first:

- **editor:** A window that enables you to edit a specific type of file or other resources. The Eclipse toolbar and menus provide different options based on the editor that is currently active.
- **feature:** A packaged set of one or more plug-ins that make up a product.
- **perspective:** A set of windows within the Eclipse GUI that are associated with a specific plug-in or task, and which are organized in a specific way.
- **plugin:** Code that performs a specific function or set of functions and can be integrated into the Eclipse framework.
- **view:** A window within the Eclipse GUI that can support an editor or provide access to a certain type of resource. Each view has a title bar that provides tabs to select different resources that are open in that view, and also often provide their own menus and icons, which are available by right-clicking or selecting an icon in the title bar.
- **workbench:** One of the two key plug-ins provided with the Eclipse IDE, the workbench plug-in provides the basic user interface, menu structure, event notification mechanism, and mechanisms for creating new windows and views. This plug-in also provides the extension points that developers can use to create menus, menu items, events, editors, and views for their plug-ins. From the user's point of view, the workbench is simply the Eclipse user interface.
- **workspace:** The second of the two key plug-ins provided by the Eclipse IDE, the workspace plug-in enables you to interact with external resources such as project directories, files, and so on. The term workspace may also be used to describe a directory used by Eclipse to hold projects.

These are my interpretations of official "Eclipse-speak." The online help for Eclipse provides the official definitions for these terms in its Workbench User Guide section.

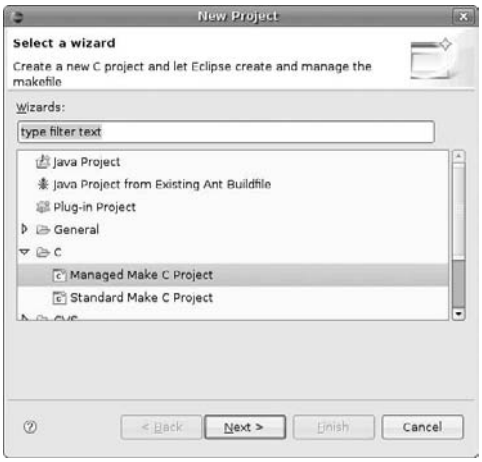
Creating a New Source Project in Eclipse

After starting Eclipse, you create new projects or project components by selecting File ➪ New, and selecting the appropriate item from this menu. To create and work with an application in Eclipse, you must first create a project to hold your source code. Eclipse provides wizards to simplify creating projects of different types. Wizards are sets of graphical screens that walk you through the process of creating a specific type of project, making sure that you don't miss any of the steps in the process.

Selecting File ➪ New ➪ Project displays the wizard selection dialog shown in Figure B3-4.

FIGURE B3-4

Selecting a Project Creation Wizard in Eclipse



Because I'm creating a C source project in this section, I've double-clicked on the C entry to see the wizards available for creating C projects, and have selected the Managed Make C Project wizard. Eclipse supports two different types of C projects—a Managed Make project in which Eclipse will automatically generate and maintain the project's Makefile for you, and a Standard Make project in which you must create and maintain your own Makefile. I'll use a Managed Make project in this example—the example in the next section shows how using a Standard Make project can simplify importing and continuing to work with existing source projects in Eclipse.

After selecting Managed Make C Project, click Next to proceed to the wizard for this particular type of project. The dialog shown in Figure B3-5 appears.

Enter the name that you want to associate with this project ("hello" in this example). You can then click Finish to create a project with default attributes (which is configured to produce a standalone Linux executable) or click Next to specify that your project produces some other type of object code, such as a shared or static library for Linux. Because this is a simple example, I'll just click Finish. The wizard creates my new "hello" project and displays a dialog asking if I want to display my new project in the C/C++ Perspective that is associated with C projects. Click Yes to display the default structure and content of the project in Eclipse's C/C++ perspective, as shown in Figure B3-6, in which I've expanded the contents of the project to highlight its default contents.

The default views and editors displayed in Eclipse's C/C++ Perspective are the following:

- **Navigator:** The view at the left, which enables you to navigate through and work with your project and the files and other resources that it contains.
- **Editor:** The pane at the top center, which differs based on the open resources in the Navigator view.
- **Console:** The bottom-right view, in which messages about interactions with Eclipse are displayed. This view also provides a Problems tab that identifies any errors or warning associ-

ated with the current project, and a Properties tab that enables you to examine any properties associated with the current project.

- **Outline:** The view at the right, which shows an outline of the files and functions referenced and defined in the active project.

FIGURE B3-5

Specifying the name of your project in Eclipse

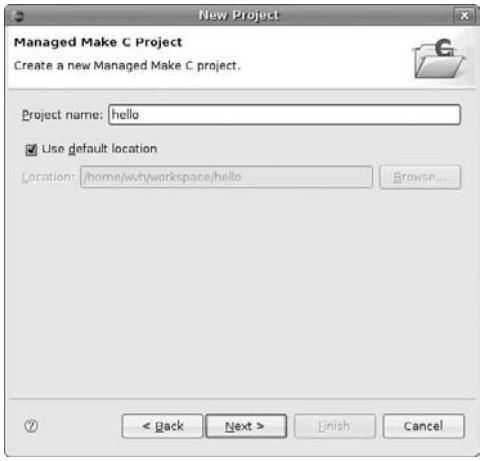
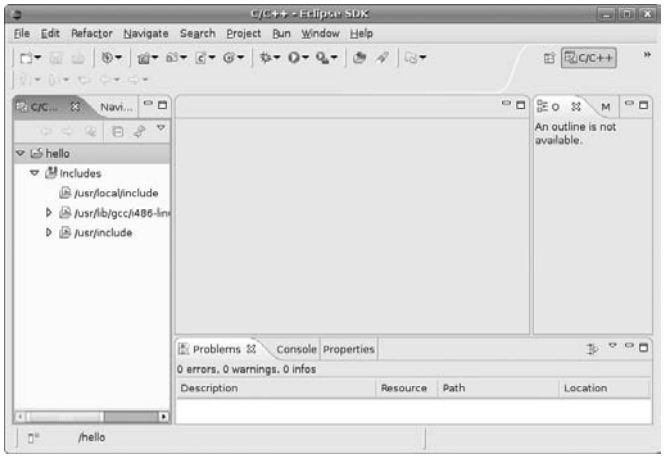


FIGURE B3-6

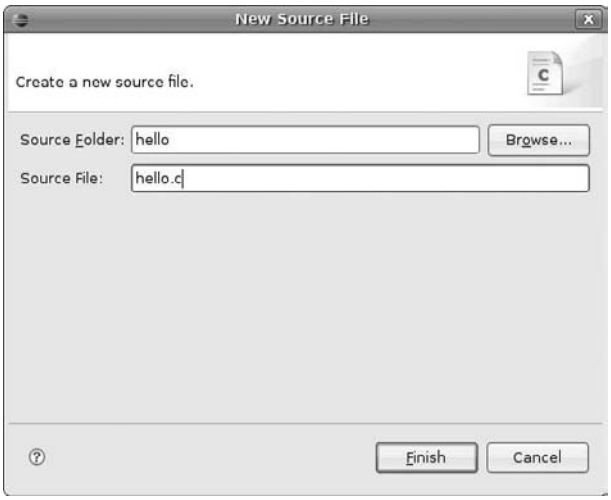
A newly created project in Eclipse



One thing missing from my sample C project is any C files, so I'll use the File ➪ New ➪ Source File menu command to create one, which displays the dialog shown in Figure B3-7. You can also create new resources for a project by right-clicking on the name of the project and selecting the type of item that you want to create from the pop-up menu's New submenu.

FIGURE B3-7

Creating a new C source file in Eclipse



Creating a new source file in a project opens the new file in the type of editor associated with that source file, as shown in Figure B3-8.

When you create your first source file in a Managed Make C Project, you'll notice that the Console view's Problems tab displays (located below the editor pane), which points out the fact that the project (which only contains one empty file) does not yet provide a main routine. Well, I suppose that's helpful—this message displays because Eclipse tries to automatically build your project each time that you change it. You can prevent this level of "helpfulness" by deselecting the Build Automatically item on the Project menu. After deselecting this option, Eclipse will only build the project when I tell it to, which is fine with me.

Next, type the source code for this module in the editor pane. Eclipse provides many keyboard shortcuts to simplify moving around in its editors. You can see the list of current key bindings for the workbench by selecting Window ➪ Preferences, and selecting General ➪ Keys in the dialog that displays, as shown in Figure B3-9.

If your fingers have already been trained to use a specific set of editor commands (as mine have been trained by the Emacs text editor), you can select different sets of key bindings by clicking the Modify tab in the dialog shown in Figure B3-9, which displays a dialog that enables you to customize your key bindings. From this tab, you can select reasonable (i.e., Emacs) key bindings by clicking the Scheme selector and selecting Emacs from the drop-down menu, as shown in Figure B3-10. You may not want to do this if you're not an Emacs aficionado, but it works for me. Click OK to close this dialog and return to the C/C++ perspective.

FIGURE B3-8

A new source file in an editor in Eclipse

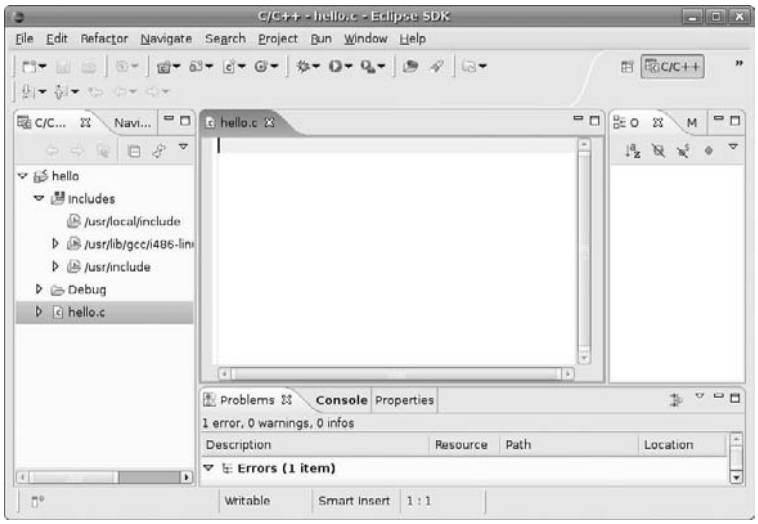


FIGURE B3-9

Viewing current key bindings in Eclipse

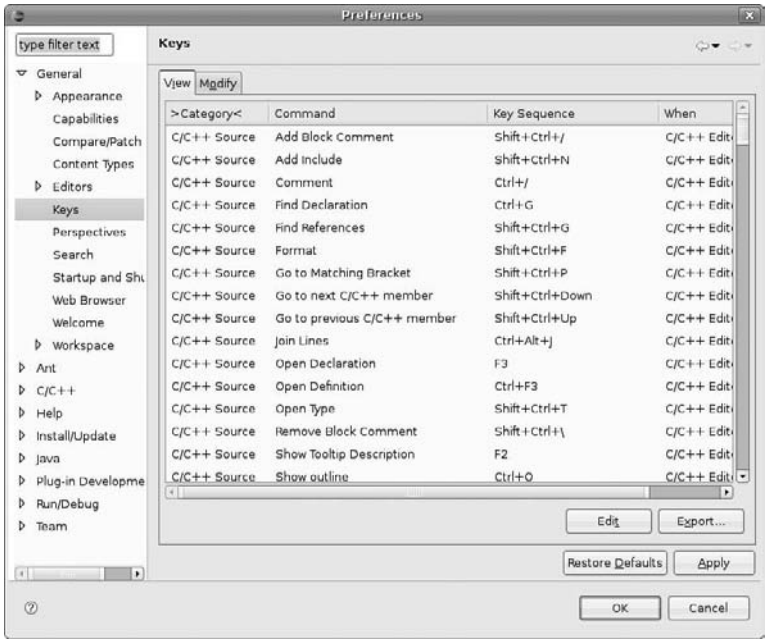
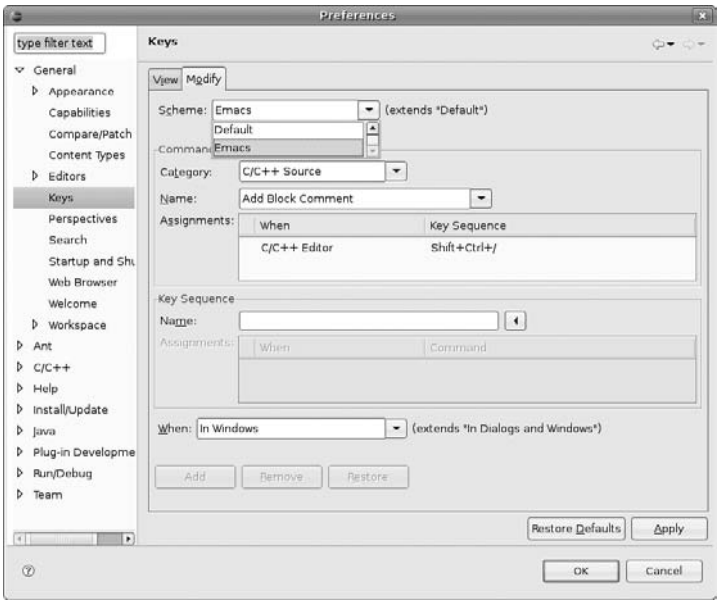


FIGURE B3-10

Customizing key bindings in Eclipse



Now that you know what commands to use to move around in the editor, type in the source code for your application (the familiar "Hello, World!" application, in this example). As you enter your source code, you'll notice that Eclipse automatically indents new lines, matches brackets, parentheses, and does the sort of code highlighting that you'd expect in a modern IDE. As you enter your code, you'll also notice that the Outline view to the left of the editor updates to identify any external references in your code, any functions that you have defined (including "main"), and so on. Figure B3-11 shows Eclipse after entering a complete "Hello, World!" program.

At this point, select File ➪ Save to save your application. As part of the save process, Eclipse analyzes your code and identifies any possible problems that it sees. Common problems are missing newlines at end-of-file, missing return values, and so on.

Now it's time to build the project. If I hadn't deselected the Project ➪ Build Automatically menu item, Eclipse would have built the project for me automatically after operations such as a Save operation. As it is, I have to build the project manually, by selecting Project ➪ Build Project (or right-clicking on the project and selecting Build Project from the pop-up menu). If you're lucky, you will not see any error or warning messages in the Console view's Problems tab, located below the editor.

The final step in this example is to actually run the application and see its output. To do this, select one of the Run commands from the Run menu. Eclipse enables you to define multiple Build and Run configurations, enabling you (for example) to define sets of compilation and execution parameters for debugging your application, executing it normally, and so on. (The default Build configuration is to build for debugging, with symbols present, and so on.) In this case, I'll select the default way of running the application, which is as a standalone, local C/C++ application by selecting Run ➪ Run As ➪ Run Local C/C++ Application, as shown in Figure B3-12.

FIGURE B3-11

A sample application and its outline in Eclipse

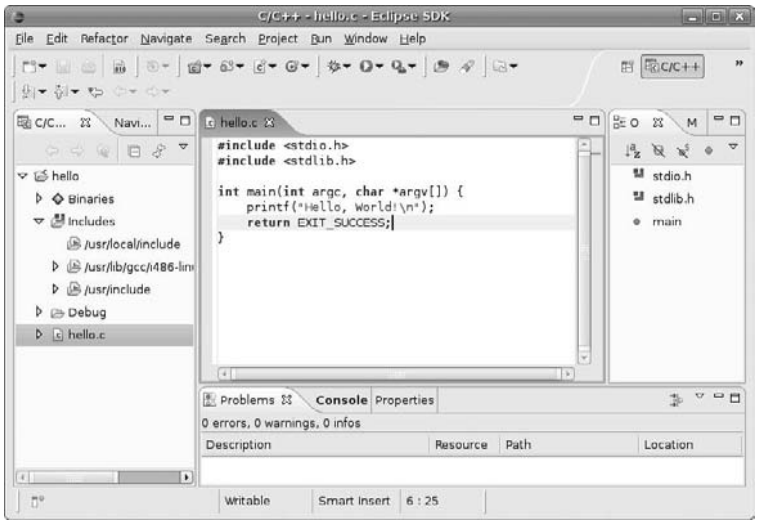
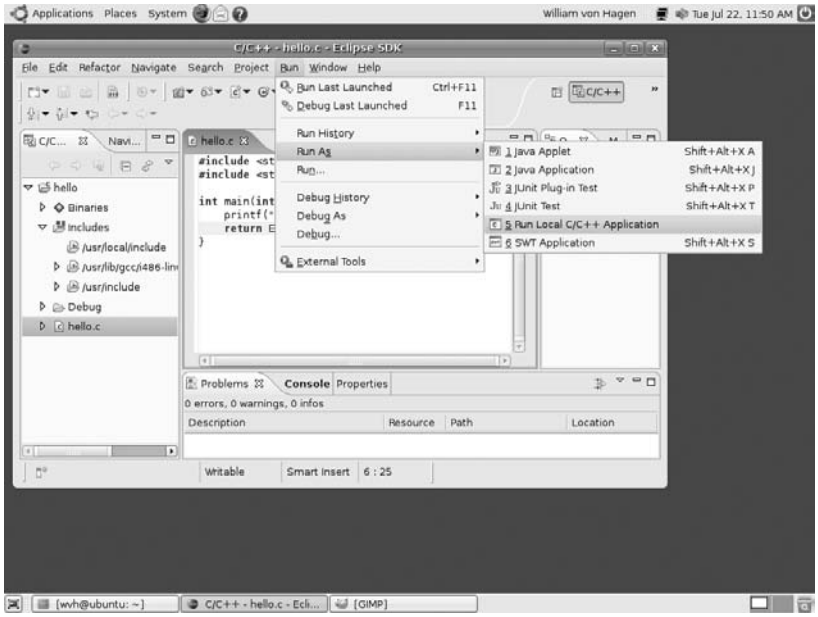


FIGURE B3-12

Running a project binary in Eclipse



Selecting this command displays the dialog shown in Figure B3-13, which enables you to select how to run this program. Because the default Build configuration used in Eclipse projects is a Debugging configuration, this dialog enables you to select whether you want to run this in the context of a local Gnu debugger (`gdb`) or a `gdb` server so that you could connect to it remotely.

Select GDB Debugger and click OK. Eclipse runs the application and refreshes the Console view, displaying the Console tab in which the output from the sample application displays, as shown in Figure B3-14.

FIGURE B3-13

Selecting a run configuration for debugging in Eclipse

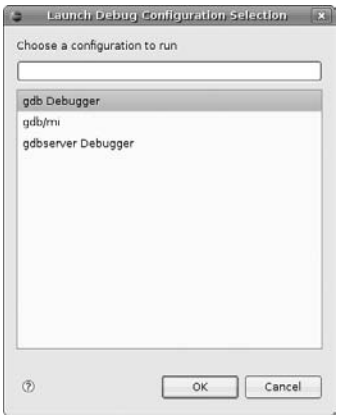
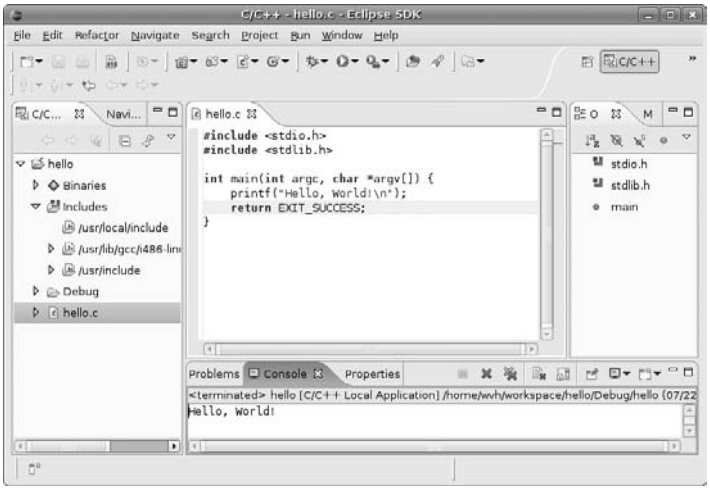


FIGURE B3-14

Running the sample application in Eclipse



To simplify running your sample application, you could create special Run configurations that would execute the application in Eclipse's other default Build configuration, which is called Release and enables you to build the application without debugging symbols, and so on, which is somewhat outside the scope of this simple tutorial. You've created a project, a source module, watched Eclipse do its magic, and even executed the application. Not bad for 10 minutes or so.

Importing an Existing Source Project into Eclipse

Creating new projects in Eclipse is one thing, but you can also take advantage of the capabilities of Eclipse in existing projects by importing them into Eclipse. This example shows how to do just that, using one of my favorite open source applications as an example, which is the `pop3check` application that enables you to list the headings for any mail that is pending for you on a remote POP mail server. I often use this as a quick, command-line way of seeing if anything interesting has come in without risking being sucked into the black hole of actually reading all of my incoming mail. The `pop3check` application was written by Steven Radack and Simon Liddington, and is available from `http://sourceforge.net/projects/pop3check`.

In my opinion, the easiest way to import an existing project into Eclipse is as a tarball that contains the source. This enables you to import a pristine version of the source code and also provides a convenient way for you to keep track of exactly what was in the project when you moved it into Eclipse, by archiving the tarball. Preparatory to beginning the import process, `cd` to the top-level of your source project directory, clean the project to remove object code, backup files, and any other items that you don't actually want to import. You'll want to generate a source tarball for your project from the top level of its source code directory so that the files that comprise your project are imported into the top level of your Eclipse project rather than being created in a subdirectory, which would happen if you imported a vanilla tarball that you retrieved from the Web or created as a backup of your project. In my case, I prepare the tarball with commands like the following:

```
$ cd pop3check-0.100.3
$ make clean
$ tar czf ../pop3check.tgz *
```

Tarball in hand, start Eclipse using the Applications ➪ Programming ➪ Eclipse command, and create an Eclipse project that corresponds to the source project that you are importing by using File ➪ New ➪ Project wizard selector. Because you're importing an existing project that already includes a Makefile, I typically select Standard Make C Project, as shown in Figure B3-15.

After specifying the name of the project and clicking Finish, the new project is displayed in Eclipse, as shown in Figure B3-16.

Because the project you've just created is a standard make project rather than one in which Eclipse will generate and manage the make information, you'll notice that the project is initially empty. Next, to import the tarball you prepared earlier, you right-click on the name of the project and select Import. This displays the dialog shown in Figure B3-17.

Because you're importing from a tarball, select Archive File and click Next. The dialog shown in Figure B3-18 displays, in which you've browsed for and selected the tarball that you prepared earlier.

FIGURE B3-15

Creating a standard C make project in Eclipse

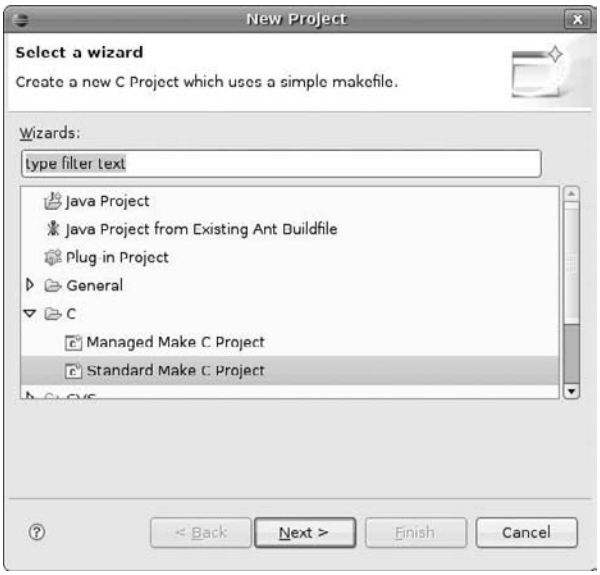


FIGURE B3-16

An empty standard make project in Eclipse

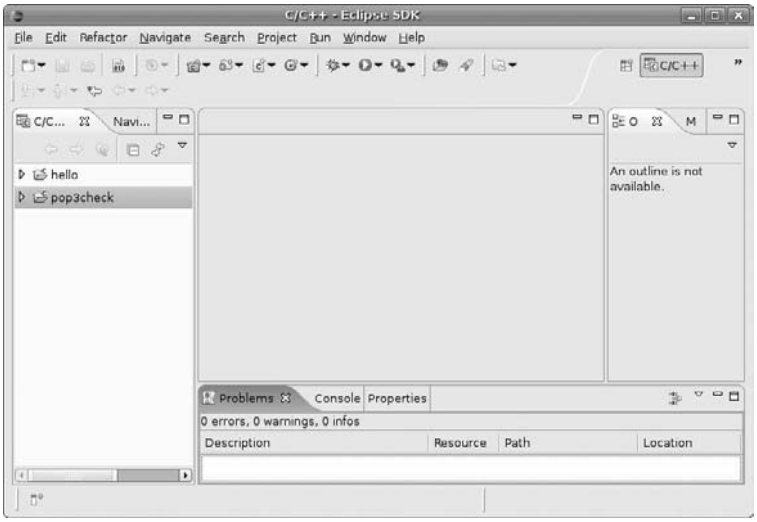


FIGURE B3-17

Specifying how to import an existing project into Eclipse

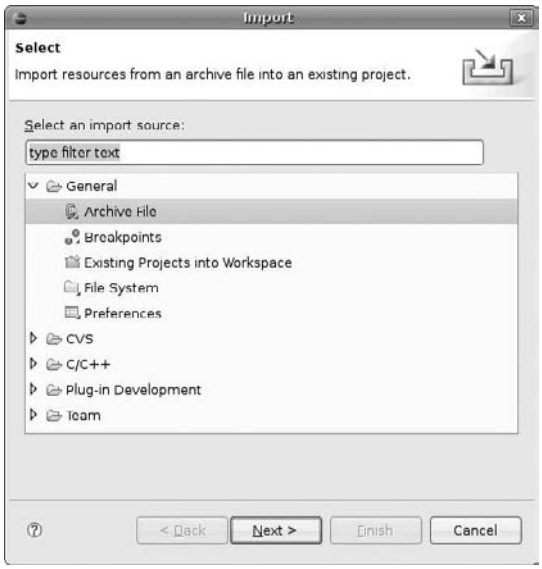
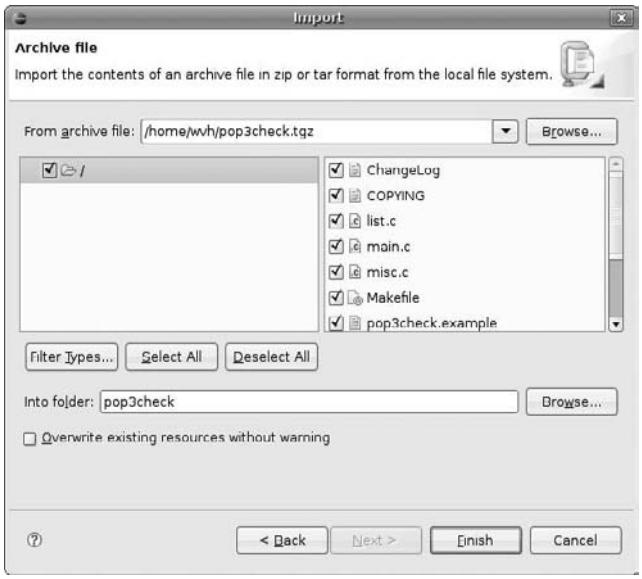


FIGURE B3-18

Selecting the tarball for importing into Eclipse



You'll note that this dialog shows an expanded view of the tarball, which enables you to verify that it contains everything you need and that Eclipse can open it successfully. Click Finish to begin the import process. Once the import completes, the contents of the imported tarball are added to the selected project, as shown in Figure B3-19.

After importing a project into Eclipse, it's a good idea to verify that the Makefile uses the build targets that Eclipse expects to find, which are all and check. To check the Makefile, double-click its name in the Navigator view, which displays the Makefile in an editor, as shown in Figure B3-20.

In this case, the default target is `pop3check`, which you could simply change to all in order to simplify things. However, since this example is an existing open source project, it's best not to change things that you don't have to in the source files. You can tell Eclipse about other build targets by right-clicking on the Makefile and selecting Create Make Target, which displays a dialog like the one shown in Figure B3-21.

This dialog enables you to define a new build target, which in this case is simply an alias for the original pop3check target. This information is saved in the Eclipse project, not in any updates to the Makefile. You can then build the project by right-clicking in the Makefile window, and selecting Build Make Target from the pop-up menu to display the dialog shown in Figure B3-22.

In this case, your newly imported project built correctly and successfully, as shown in Figure B3-23.

FIGURE B3-19

An imported project in Eclipse

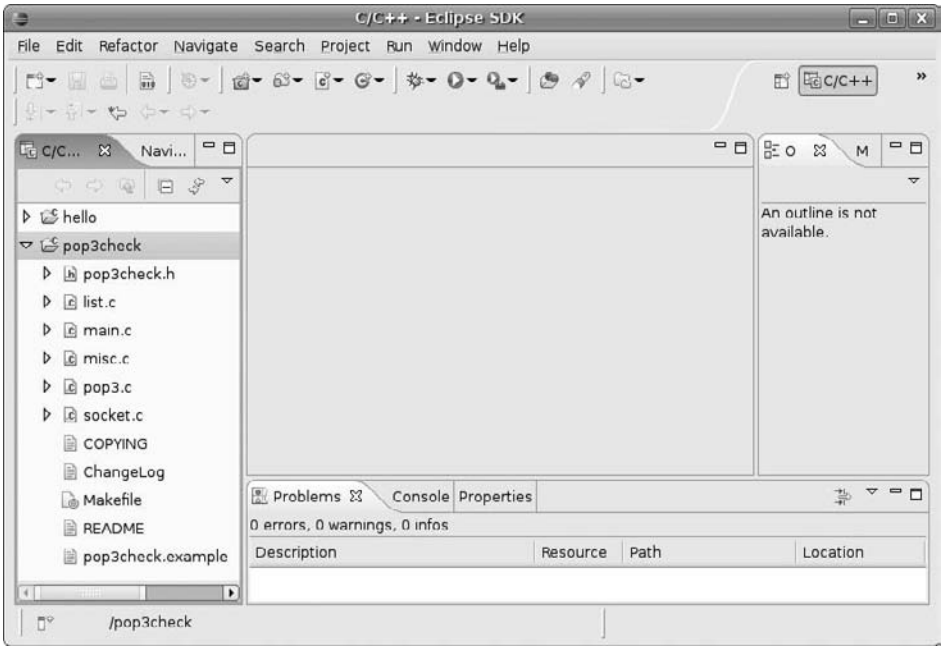


FIGURE B3-20

Examining the Makefile for an imported make project

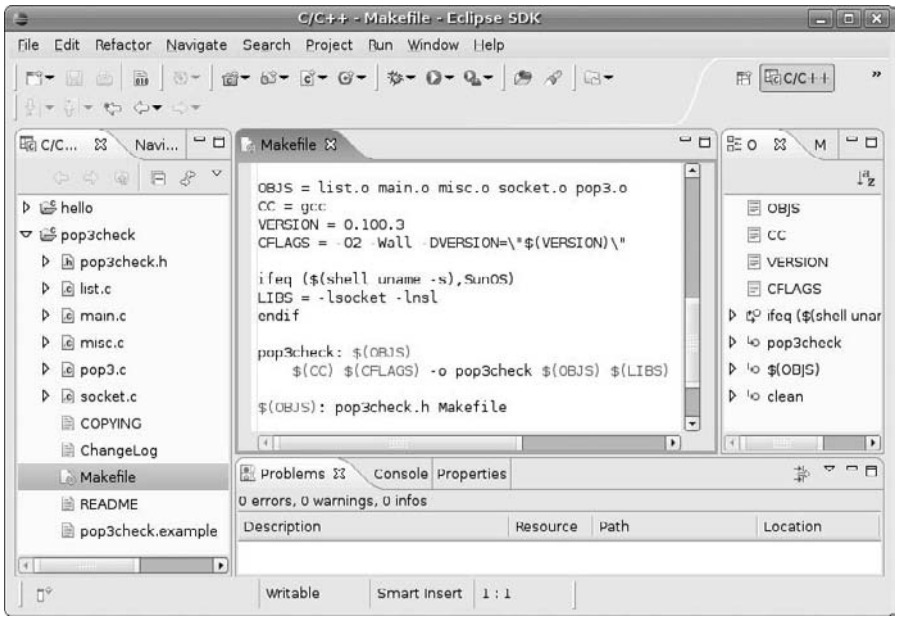


FIGURE B3-21

Defining a custom Make Target for an imported project

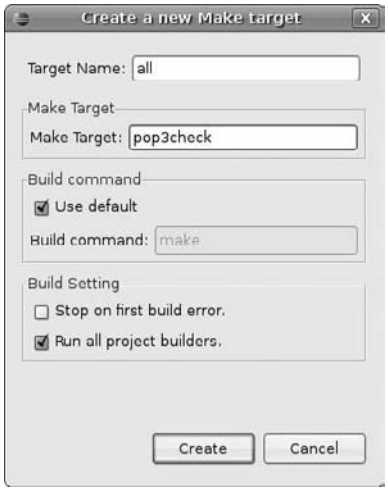


FIGURE B3-22

EBuilding a custom Make Target in Eclipse

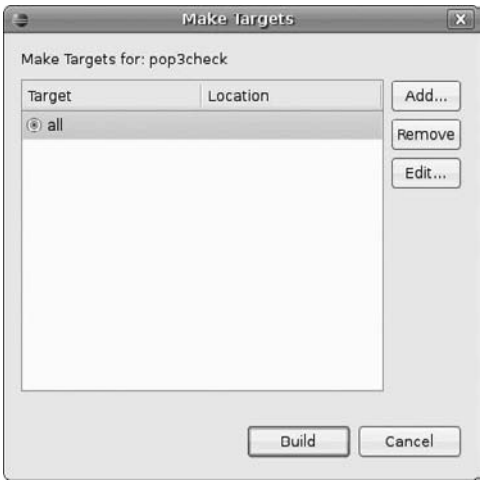
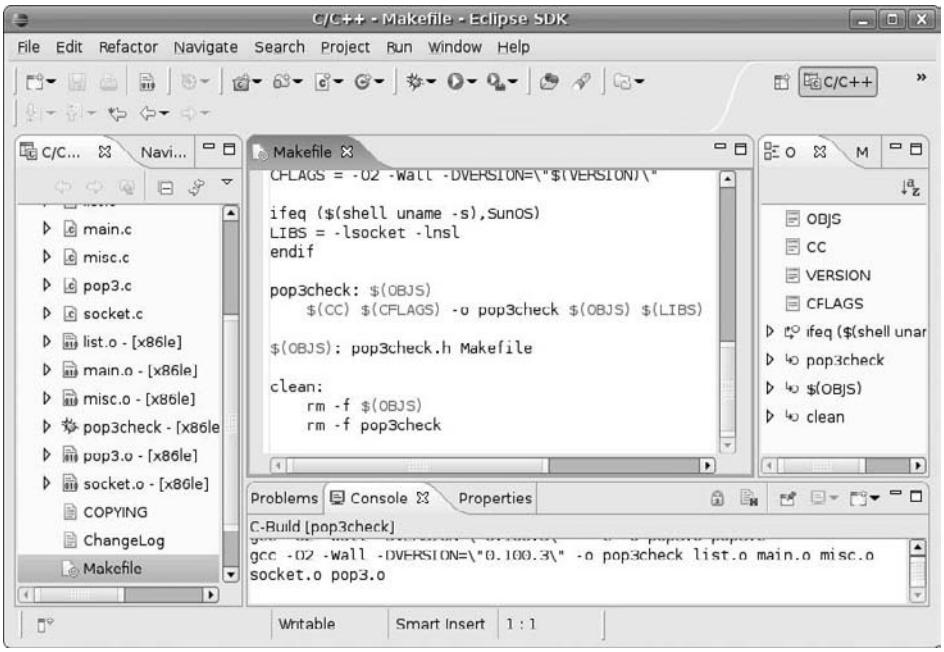


FIGURE B3-23

Console output from a build of an imported project in Eclipse



As you might expect, your mileage may vary when importing existing source projects into Eclipse, but the ease with which you can do so makes it easy for you to move your existing application development efforts into Eclipse so that you can take advantage of its many capabilities and the ease of use that a modern IDE such as Eclipse provides.

Getting More Information about Eclipse

Eclipse is a powerful, full-featured IDE, which means that the examples in the previous two sections can only scratch the surface of its capabilities and bells and whistles. For more detailed information about using Eclipse, see one of the many books on using Eclipse that are available, or visit some of the following resources on the Web:

- `www.eclipse.org`: home page of the Eclipse mothership
- `www.eclipse.org/documentation/main.html`: Eclipse documentation site
- `http://planeteclipse.org/planet`: Planet Eclipse, frequented by Eclipse developers, contributors, hackers, and users
- `www.eclipseplugincentral.com`: Eclipse Plug-in Central, from which you can download many free Eclipse plug-ins

The next section talks about source code control software, highlighting the Subversion SCCS. As an example of the wide range of free plug-ins available for Eclipse, the following are two free Subversion plug-ins for Eclipse:

- Subversive Plug-in at `www.polarion.org`
- Subclipse Plug-in at `http://subclipse.tigris.org`

Not only do these plug-ins provide an easy way to work with Subversion repositories in Eclipse, but they also provide a great example of how all of this open source stuff tends to fit together nicely. And you certainly can't beat the price!

Source Code Control Software

If you've ever accidentally deleted a source file that you were working on, lost the software version of the great American novel when your system crashed and lost its disks, or worked on a joint project with someone who was thoughtful enough to overwrite all of your new code, you might have wished that you'd been using a source code control system once you finished swearing. A source code control system, popularly known as an *SCCS*, is a software package that manages a repository of source code, tracking changes to existing code, managing different versions of that code (often known as *branches*), and providing a mechanism for recovering any version of any file in the repository at any point in time.

If I had to count the number of times that an SCCS had saved my bacon when doing development, I would quickly run out of fingers and toes. This doesn't necessarily mean that I'm sloppy. When multiple developers work on the same sets of files, conflicts between their changes are bound to occur. An SCCS provides a mechanism for identifying and resolving conflicting sets of changes (known as *merges*). Similarly, a source code control system facilitates experimentation with code by providing a safety net that you can use to fall back to versions of files that you know to be working, just in case your latest bright idea didn't turn out to be very illuminating after all.

CAUTION **One common misconception is that an SCCS is a replacement for backups. Not really true. An SCCS may be a replacement for backing up your own files, because you can check them in and retrieve them later if your machine bursts into flames. However, the SCCS repository itself must be backed up regularly—if anything, a central repository used across a project or enterprise provides a single, critical point of failure.**

The next two sections discuss the two most popular source code control systems for Linux, the Concurrent Versioning System (CVS) and the Subversion Project. I have used CVS for years, but am purposely giving it short shrift here—Subversion is often referred to as CVS++, and for good reason. Not only is Subversion a more modern implementation with a richer set of commands than those provided by CVS, but also Subversion is based on modern protocols and is more easily integrated into Web browsers, integrated development environments, and graphical computer interfaces, and is actively under development. Alas, poor CVS, I knew thee well. . . .

The Concurrent Versioning System (CVS)

My humorous comments about CVS in the previous section are meant to be just that, and should not be taken to mean that CVS is not a time-tested, stable SCCS. CVS has been in use and under development for years, and is still often used to provide access to the source code for open software projects over the Internet. That said, most projects that I'm aware of which have traditionally depended on CVS are mobbing to Subversion, and that's a trend that I believe will continue, if not accelerate. Because neither CVS nor Subversion is installed as part of a basic Ubuntu installation, I'd strongly suggest that you install and use Subversion rather than CVS. However, if you're a CVS fanatic or your current projects already work with CVS, you can use the Synaptic Package Manager to install CVS, and then see sites such as the following for information about setting up and Web-enabling CVS repositories:

- `www.faqs.org/docs/ldev/0130091154_198.htm`: The Linux Development FAQ's section on CVS.
- `www.freeos.com/articles/4608`: A nice (and short!) article by Mahantesh. M.Vantmuri on setting up and configuring CVS and synchronized copies of CVS repositories. Though this article is a few years old, not much has changed—we are talking about CVS, after all.
- `www.flexwiki.com/default.aspx/FlexWiki/ConfiguringCVS.html`: A nice wiki that provides a lot of good information about understanding and using CVS.

CVS is a fine system, and I've used it extensively at many companies. However, I feel that if you're going to set up an SCCS from scratch on your Ubuntu system, you may as well use the more modern and powerful Subversion SCCS.

Installing and Using Subversion

Subversion (`http://subversion.tigris.org`) is a relatively new, yet extremely popular source code control system (SCCS). Aside from its power and rich command set, much of Subversion's popularity lies in the fact that it is based on WebDAV (Web Distributed Authoring and Versioning), a Web-based protocol that enables you to view and access Subversion repositories in any Web browser. This includes the ability to drag and drop files from your Web browser to your desktop to check them out, and vice versa to check them in. As you'll see in the next few sections, setting up a Subversion repository and integrating it into your Web server is quite easy to do, just as Subversion is easy to use from the desktop or the command line.

Installing Subversion

You can install Subversion using any of the commands described in Chapter 23, "Adding, Removing, and Updating Software." When installing Subversion, I generally install the following packages, not all of which are required. My rationale for each follows the name of the package in this list:

- `libapache2-svn:` (Optional) This package is necessary only if you want to provide a Subversion repository that users can access through your Web server. Selecting this item causes the Apache2 Web server and related files to be installed on your system. I always find it handy to be able to access my repository over the Web, it's easy to configure everything at once, and the extra disk space that Apache2 consumes is always worth it to me. However, you can always add this in later. See Bonus Chapter 8, "Setting Up a Web Server," for information about manually installing different versions of the Apache Web server.
- `rapidsvn, esvn, or kdesvn:` (Optional) Graphical clients for Subversion. Only necessary if you want to use a graphical client. The `rapidsvn` client is quite popular and is a GNOME client. The `esvn` client uses the Qt graphical interface and, therefore, doesn't have much baggage in either the GNOME or KDE desktops. The `kdesvn` client is a KDE client and, therefore, requires that the KDE libraries be installed on your system, but it is the only Subversion client that accepts command-line arguments. The `kdesvn` client can, therefore, be handy to use if you want to integrate a graphical Subversion client into Firefox to handle `svn:` URLs. See the sidebar entitled "Handling svn: URLs in Firefox" later in this chapter, for more information about doing this.
- `subversion:` (Mandatory) This is the core Subversion package, and provides everything you need to use Subversion from the command line on your local system.

TIP **You may also find it handy to install the `subversion-tools` package when installing Subversion. This package provides a few tools for sharing Subversion repositories between users and doing more sophisticated imports of existing source code directories into Subversion. However, if you have not already installed a mail server on your system, selecting this package for installation causes the `exim` mail server and related files to be scheduled for installation on your system. I prefer the Postfix mailer, so I always install and configure a mail system before installing Subversion on a system that is also a real mail server. In other cases, where the Subversion server is separate from my mail server, I just accept that fact of some lost disk space. See Bonus Chapter 9, "Setting Up a Mail Server," for general information about the different mail servers that are available on Ubuntu.**

Once the installation completes, you're ready to set up your repository and, if you installed the `libapache2-svn` package, configure Apache2 to support Subversion.

See Chapter 23, "Adding, Removing, and Updating Software," for more detailed information on installing applications.

Setting Up a Subversion Repository

Creating a Subversion repository is done with the `svnadmin create` command, which creates the directory structure that makes up a Subversion repository. This command takes one argument, which is the name of the directory in which you want your Subversion repository to be created.

One important thing to remember when creating a Subversion repository is that, because Subversion depends on WebDAV, which depends on the HTTP protocol, your Subversion repository must be owned by the user and group that is used for communication via this protocol. On Ubuntu, this is the www-data user and the www-data group.

TIP **If you are creating a Subversion repository named `svn` and plan to export it through a Web server by that name, you must not create it in your Web server's `DocRoot` directory. As you'll see later in this section, you specify the full pathname to the Subversion repository in Apache's Subversion configuration file. Trying to use a repository named `svn` that is located in your Web server's `DocRoot` directory will result in the error `301 Moved Permanently (http://server-name)`, which isn't exactly intuitive.**

For example, to create a Subversion repository in the directory `/var/svn`, you would execute the following command:

```
$ sudo-u www-data svnadmin create /var/svn
```

After supplying your password, the `svnadmin create` command creates your repository and populates it with the subdirectories it requires, all owned by the proper user and group. You can verify this by listing the contents of the specified directory, as in the following example:

```
$ ls -l /var/svn
total 28
drwxr-xr-x 2 www-data www-data 4096 2006-05-29 12:31 conf
drwxr-xr-x 2 www-data www-data 4096 2006-05-29 12:31 dav
drwxr-sr-x 5 www-data www-data 4096 2006-05-29 12:31 db
-r--r--r-- 1 www-data www-data    2 2006-05-29 12:31 format
drwxr-xr-x 2 www-data www-data 4096 2006-05-29 12:31 hooks
drwxr-xr-x 2 www-data www-data 4096 2006-05-29 12:31 locks
-rw-r--r-- 1 www-data www-data  229 2006-05-29 12:31 README.txt
```

Of course, an empty Subversion repository isn't all that exciting. The next section explains how to add an existing project to your repository, so that you actually have some source code to control.

Importing Projects into a Subversion Repository

A Subversion repository is simply a catcher's mitt for the files and directories that make up different projects. You can simply check your existing files directly into a Subversion repository, but it's far better to check files in using a separate directory hierarchy for each project.

To check a project directory into Subversion, first remove all nonsource files such as object files, intermediate files, binaries, and so on from that directory. Next, you simply import the project into Subversion using Subversion's svn `import` command, as in the following example where you import a project that contains the canonical C program, "Hello, World!":

```
$ svn import hello file:///var/svn/hello -m "Creating hello project"
Authentication realm: <http://192.168.6.90:80> Subversion Repository
Password for 'wvh':
Adding         hello/hello.c
```

The arguments to the `svn import` command are the project that you want to import, the Subversion URL of the location in the repository where you want the project to be created, and the `-m` (message) option, followed by a string that contains the initial check-in message for the project, enclosed within quotation marks. The `-m` option is optional, of course, but it saves time because you would otherwise have to enter this message in a text editor that the `svn` command would start for you.

Once you've created a project or checked any files into Subversion, you can use the `svn` list command to display the contents of the repository, as in the following example, which uses the `-R` option to this command to recursively display the contents of the repository:

```
$ svn list -R file:///var/svn
hello/
hello/hello.c
```

You'll notice that this example uses a URL beginning with `file:///`, which provides access to a local repository. This is necessary at this point because you haven't started a Subversion server, so all URLs must be specified relative to the current filesystem. Trying to access your Subversion repository using any Web-based URL would cause an error message to be displayed at this point, either svn: PROPFIND of `'/svn': could not connect to server` (if you do not have a Web server running) or svn: `PROPFIND of '/svn': 405 Method Not Allowed` (if you have a Web server running, but no Subversion repository is available). The next section explains the two Subversion servers that are provided for use on Ubuntu systems and how to configure them to make your repository available from other machines.

Setting Up a Subversion Server

Although Subversion is based on WebDAV and, therefore, uses the HTTP protocol for basic communications between clients and servers, you can access a local Subversion repository (i.e., one that is physically located on the same machine that you are using) without running a Subversion server. However, you will need to be running a Subversion server on machines whose repositories you want to access from other systems. There are two common types of Subversion servers:

- **Apache2:** An Apache Web server with integrated support for Subversion. Support for Subversion and WebDAV is easily added to an Apache or Apache2 Web server, which enables you to access your Subversion repository over the Web using URLs like `http://host-name/repository or https://host-name/repository` (if you configure SSL support into Apache2 and Subversion). The examples in this chapter discuss Apache2 integration and configuration, because Apache2 is the more modern and powerful Apache Web server.
- **Svnserve:** A standalone Subversion server that enables you to access your Subversion repository over the Web using URLs like `svn://host-name/repository. The svnserve` binary is installed as part of the Subversion package.

These two servers are not mutually exclusive and can easily run on the same machine at the same time, providing access to the same repository if correctly configured. The next two sections explain how to configure each of these to provide access to the sample repository that you set up and populated in previous sections.

Configuring Apache2 for Subversion Support

Configuring the Apache2 Web server to support Subversion is quite easy, only requiring that you install and integrate the right Apache module, configure that module, and create the authentication file that will control write access to the repository.

Installing the Apache2 Web server's support for Subversion was described earlier in this chapter in the section entitled "Installing Subversion," and is done by installing the `libapache2-svn` package using the Synaptic Package Manager or the `apt-get utility`. The Apache2 Web server's support for Subversion

integration with Apache2 is loaded by the file `/etc/apache2/mods-available/dav_svn.load`. The configuration file for this module is the file `/etc/apache2/mods-available/dav_svn.conf`.

After installing the Apache2 Web server's support for Subversion, edit this file using `sudo` and your favorite text editor. Your configuration options are the following:

1. Update the `<Location /svn>` element if you want to make your repository available through any URL other than `http://Web-server/svn`.
2. Uncomment the `DAV svn` entry to enable the repository.
3. Edit the `SVNPath` entry to identify the full pathname of your Subversion repository in the server's directory structure.
4. To use basic HTTP authentication (which I will use in this example), uncomment the following three lines:

```
AuthType Basic
AuthName "Subversion Repository"
AuthUserFile /etc/apache2/dav_svn.passwd
```

5. Modify the pathname in the `AuthUserFile` entry to identify the HTTP password file that you want to use for HTTP authentication.
6. Purge any existing HTTP password file for Subversion and create the HTTP password file using the following commands, where `user-name` is the name of the first user for whom you want to create a password entry:

```
$ sudo rm -f /etc/apache2/dav_svn.passwd
$ sudo htpasswd2 -c /etc/apache2/dav_svn.passwd user-name
  New password:
  Re-type new password:
```

7. Add any other users to the HTTP password file using the following command:

```
$ sudo htpasswd2 /etc/apache2/dav_svn.passwd user-name
```

After editing this file, a sample `dav_svn.conf` file looks like the following example. I've removed the comments to make things easier to read:

```
<Location /svn>
  DAV svn
  SVNPath /var/svn
  AuthType Basic
  AuthName "Subversion Repository"
  AuthUserFile /etc/apache2/dav_svn.passwd
  <LimitExcept GET PROPFIND OPTIONS REPORT>
    Require valid-user
  </LimitExcept>
</Location>
```

At this point, all you have to do is restart the Apache2 Web server, which you can do using the `sudo /etc/init.d/apache2 restart` command, and your repository will be available from other machines through the URL `http://Web-server/svn`. See "Checking Files out of a Subversion Repository" later in this section, for information about using this URL to check out the files for your projects.

Configuring the svnserve Server

As mentioned earlier, Subversion's `svnserve` server enables Subversion clients to use URLs that use the `svn` protocol directly rather than the HTTP protocol. The `svnserve` server is installed as a part of the basic Subversion package.

You can start the `svnserve` server manually or through an Internet services daemon such as `inetd` or `xinetd`. The `svnserve` server doesn't require any special configuration because you identify the root of your Subversion repository on the command line or in the configuration files for any Internet services daemon that you want to use. To start the `svnserve` daemon from the command line to provide access to the repository you created earlier in this chapter, you would simply execute the following command:

```
$ sudo -u www-data /usr/bin/svnserve -r /var -d
```

This starts the `svnserve` server as the user www-data, uses the -r option to identify the root of the Subversion repository as /var, and uses the -d option to start the svnserve process as a daemon that runs in the background.

As mentioned previously, you can also configure an Internet services daemon to handle incoming requests for `svn:` URLs on demand, rather than starting the `svnserve` daemon manually each time that you log in or explicitly add it to your Ubuntu system's startup files. Internet services daemons wait for incoming requests for services on certain ports and start associated daemons on demand, rather than having the daemons always running. Ubuntu provides two Internet services daemons, `inetd` and `xinetd`. The `inetd` daemon is the original Internet services daemon, while the `xinetd` daemon is a more modern and flexible Internet services daemon. Neither of these is installed by default on a basic Ubuntu system, so you will have to manually install whichever one you choose (I suggest `xinetd`) using the Synaptic Package Manager or apt-get.

If you install the `xinetd` Internet services daemon, you will also have to create the file /etc/xinetd.d/svn, which is the configuration file that tells the `xinetd` daemon what to do when an incoming `svn` request is received on port 3690, as defined in the file /etc/services. An /etc/xinetd.d/svn file that provides access to the repository that you created earlier in this chapter would look like the following:

```
service svn
{
    disable = no
    socket_type = stream
    protocol = tcp
    wait = no
    user = www-data
    server = /usr/bin/svnserve
    server_args = -i -r /var
}
```

After creating this file, you must restart the `xinetd` daemon using a command like the following:

```
$ sudo /etc/init.d/xinetd restart
```

Handling svn: URLs in Firefox

Trying to access a URL that begins with `svn:` in Firefox will display a dialog with the message "Firefox doesn't know how to handle this address because the protocol (svn) isn't associated with any program." If you want to be able to handle these URLs in Firefox, you'll need to create a small wrapper script that invokes your favorite graphical Subversion client with that URL as an argument. For example, I use the following script to automatically start the `kdesvn` client (the only Subversion client that accepts command-line arguments at the moment) when I click on an svn: URL:

```
!#/bin/bash
/usr/bin/kdesvn "$1"
```

After saving this script in a file and making it executable, I modified Firefox to call it when I click on an `svn:` URL by doing the following:

1. Type `about:config` in the Firefox location bar, and press Return.
2. Right-click in the Preference Name window and select New ➪ String.
3. Enter `network.protocol-handler.app.svn` as the name of the string. Click OK.
4. Enter the full pathname of the script as the value of the string.

This isn't especially exciting, but it does prevent occasional frustration when clicking on `svn:` URLs. It would be nice if `esvn` or `rapidsvn` took command-line arguments, because `kdesvn` is a KDE application and installs a ton of KDE libraries if you haven't already installed other KDE applications on your Ubuntu system. One of these days . . .

Checking Files out of a Subversion Repository

Now that you've configured one (or two) Subversion servers, you can access the sample repository from any machine that can reach your server. Sample checkout commands that would retrieve the entire hello project are the following:

- `svn co file:///var/svn/hello:` This is the standard command to retrieve the hello project. You can use this command only from the machine on which the Subversion repository is located.
- `svn co http://192.168.6.90/svn/hello:` This command retrieves the hello project using the HTTP protocol, and it will work only if you have a Web server running on the host where the Subversion repository is located, as described in the section entitled "Configuring Apache2 for Subversion Support."
- `svn co svn://192.168.6.90/svn/hello:` This command retrieves the hello project using the SVN protocol, and it will work only if you have the Subversion server running on the host where the Subversion repository is located, as described in the section entitled "Configuring the svnserve Server."

Executing any of these commands will create a directory named `hello` in your current working directory and will populate that directory with the file `hello.c`. Congratulations—you're using Subversion!

Committing Changes to a Subversion Repository

Once you have checked files out of your Subversion repository and modified them, you'll want to check them back into the repository from which they came—providing one true location for the latest and greatest version of your source code and being able to track changes to it is the whole point of source code control in the first place.

Subversion makes it very easy to check modified files back in because it automatically keeps track of version information and the repository from which checked-out files came in the first place. If you do a long listing (`ls -al`) of any directory that you've checked out of Subversion, you will see a directory named `.svn` that holds this information.

Because Subversion maintains this information for you, all you have to do to check a file back into Subversion is to use the `svn ci` (check in) command, as in the following example, which checks in a new, improved version of the file `hello.c` and shows the output from the `svn` command's check-in operation:

```
$ svn ci hello.c -m "Added comment"
Sending        hello.c
Transmitting file data .
Committed revision 7.
```

The `-m` option enables you to specify your check-in message on the command line. If you do not use this option, the `svn` command will automatically start a text editor in which you can enter a check-in message and will proceed with the check-in once you exit the editor.

Getting More Information about Subversion

Subversion is a powerful source code control system, with many more commands and options than I've needed to discuss in the previous few sections. Some great Web sites for additional information about Subversion are the following:

- **Subversion Book** (`http://svnbook.red-bean.com`): An online version of the O'Reilly book *Version Control with Subversion*. Though this is available online for your convenience, it's always nice to buy a copy (http://svnbook.red-bean.com/buy/p), because we all understand that authors like to be compensated for their work other than through "good vibes" and "thank you" messages.
- **Subversion FAQ** (`http://subversion.tigris.org/faq.html`): The Subversion FAQ, which is a great source of answers to common questions and general information about Subversion.
- **Subversion Home Page** (`http://subversion.tigris.org`): This is the home page for Subversion, with as much information as you'd expect from the Subversion mothership.
- **Subversion IRC:** If you need answers *now* and are a fan of Internet Relay Chat, a dedicated channel for Subversion users is available at `irc.freenode.net`, channel #svn.
- **Subversion Mailing Lists:** The tigris.org folks provide a variety of mailing lists for Subversion information. The `announce@subversion.tigris.org` mailing list is a low-traffic, read-only list to which release and update announcements are posted. The `dev@subversion.tigris.org` list is a mailing list for Subversion developers. The `users@subversion.tigris.org` list is a high-traffic list for Subversion users where you can ask questions and help out other Subversion users. To subscribe to any of these lists, send mail to `listname-subscribe@subversion.tigris.org` and follow the instructions that you receive in the response mail that you receive. These mailing lists are moderated, so there may be a delay before any posts that you make actually appear on the lists.

Many other Subversion-related pages and sites are available on the Web; a simple Google or Clusty search for Subversion will keep you knee deep in reading material for weeks. The Indian *Linux for You* magazine (`www.linuxforu.com`) published a nice article on Subversion, which I couldn't find in the archives on their site, but is available from the authors' site at `http://mia.ece.uic.edu/~papers/publications/subversion_final.pdf`.

Summary

You don't have to build your own Linux applications anymore unless you want to—providing a precompiled, integrated set of Linux applications, libraries, and a kernel is the whole point of a Linux distribution. However, you may sometimes want to build your own, bleeding-edge versions of existing software to get the latest and greatest features, or you may be a developer and want to write some of your own tools (hopefully open source) because that's what you do. This chapter explained how to install the tools that you'll need to compile most C and C++ applications on an Ubuntu system, including related software such as the `make` build tool and various libraries. The chapter also explored many of the most popular IDEs available for Linux systems, primarily focusing on the popular, powerful, and open source Eclipse IDE. The chapter concluded by discussing source code control systems and how to set up your own Subversion repository and server, and how to quickly and easily make your Subversion repository available through a local Apache Web server.

Bonus Chapter 4 continues with system-related topics, discussing how to create and manage users and groups on your Ubuntu system, and why this is important. It also discusses advanced topics related to file and directory protection, such as Access Control Lists, that are above and beyond the traditional user/group protection model but give you even finer control over who can access your data (and don't require system administration privileges to do so).

Bonus Chapter 4

Managing Users, Groups, Authentication, and Advanced Permissions

Although many people install and use Ubuntu as the operating system on a truly personal computer, Linux is designed and implemented as a multiuser computer system from the group up. Even if you're the only user of an Ubuntu system, that system still has multiple user and group accounts and runs many processes using the rights granted to those accounts. Similarly, the login process seems simple, but actually invokes several user and group checks under the hood, as do other commands that require specific permissions.

This chapter explains how to manage users and groups on Ubuntu and Kubuntu systems. Chapter 4, "Basic Linux System Concepts," provided a basic introduction to the general Linux concepts of users, groups, file permissions, and performing privileged operations. The following is a quick review:

- The section of Chapter 4 entitled "Basic Concepts: Users and Groups" provided an overview of how users and groups work on Ubuntu and other UNIX-like systems, explaining the content and organization of the system files that hold user and group information.
- The section of Chapter 4 entitled "File and Directory Permissions under Linux" introduced file and directory permissions and explained what the cryptic protection entries in a long file listing actually mean.
- The section of Chapter 4, "Performing Privileged Operations in Ubuntu," explained how Ubuntu handles the execution of privileged commands, which is a slight departure from how other Linux and UNIX-like systems work.

This chapter delves into related topics in more detail. The information in this chapter ranges from a hands-on discussion of how to create and modify user and group accounts using graphical tools to more detailed discussions of how users and groups actually interact with your Ubuntu or Kubutnu system and the files and directories that it contains. The first section discusses the graphical and command-line tools used to create and manage user and group

IN THIS CHAPTER

Creating users and groups

Understanding standard Linux authentication

Configuring the sudo command

Using ACLs from the command line

Editing ACLs graphically

Installing and Using Kerberos

accounts. The following section explains how authentication works when a user logs in, which should be of interest to Ubuntu and Kubuntu system administrators in enterprise or academic environments that may require changes to the standard authentication process to support custom authentication mechanisms. This chapter also discusses the details of how Ubuntu and Kubuntu systems handle the execution of privileged commands, which is a slight departure from how other Linux and UNIX-like systems work. The next section in this chapter discusses Access Control Lists (ACLs), which provide fine-grained control over who can access selected files, and includes an ACL-to-earth dictionary to understand the cryptic ACL commands, their obscure output, and their motto, which seems to be, "Is that gibberish on the screen, or are you just glad to see me?" The final section in this chapter discusses how to install and use Kerberos, a secure network-oriented authentication mechanism that is used in many large enterprise and academic environments because of the high level of security that it provides.

NOTE **Ironically, Kerberos, an open source project, is also the authentication mechanism that underlies authentication in Microsoft's most secure environments, such as Active Directory, Windows Server 2003, Windows Server 2008, and so on.**

Creating and Managing Users and Groups

Users and groups are primarily ways of granting different permissions to different individuals or sets of individuals, respectively. These permissions include the ability to read, write, and access specific files and directories; the ability to access specific devices in different ways; and the ability to read or write to and from special-purpose resources such as filesystems, portions of memory, and kernel resources.

TIP **The following sections focus on the graphical tools for user and group creation that are provided on Ubuntu and Kubuntu systems. Command-line tools for quick user and group creation, modification, and removal are also available on all Ubuntu and Kubuntu systems. For more information about these tools, see the online reference information for the `adduser`, `useradd`, `userdel`, `usermod`, `groupadd`, `groupdel`, and `groupmod` commands. I prefer to use the graphical tools when creating single users to take advantage of capabilities such as their default privilege settings and associated group memberships. The command-line tools are quite useful in scripts, but they are also well documented, and I don't see much point in simply recreating that documentation here. One tip, though—when you're using these commands to add users to existing groups, either use the `adduser` command (`adduser user-name group-name`) or make sure that you specify the `-a` option when using the `usermod` command (`usermod -a -G group-name user-name`). Forgetting to use the `-a` option with `usermod` will make the specified `user-name` a member of `group-name` and only that group—it will remove `user-name` from all other groups. I can personally attest to the fact that it is painful to remember and add back all of the other group memberships that most desktop users require.**

Managing Users and Groups on Ubuntu Systems

Ubuntu provides a convenient graphical tool for creating users, creating groups, and adding users to groups. Select the System ➪ Administration ➪ Users and Groups menu entry to execute this tool. The dialog shown in Figure B4-1 appears.

In order to do anything other than simply see the lists of users and groups in this tool and view their properties, you must click Unlock, enter your password in the dialog that appears (or select another user with administrative privileges and enter their password), and click Authenticate to proceed. This enables the Add User button on the Users Settings dialog and the Add Group button on the Groups Settings dialog, and also enables you to edit the properties of any existing user or group.

FIGURE B4-1

Ubuntu's Users and Groups tool



TIP **For information about using the Users and Groups tool beyond what is discussed in the next few sections, see its online manual, which you can access at any time by clicking Help on any panel of the main Users and Groups dialog.**

Creating New Users

To add a new user using the Users and Groups tool, click Add User to display the New User Account dialog's Account tab, shown in Figure B4-2.

FIGURE B4-2

The New User Account dialog

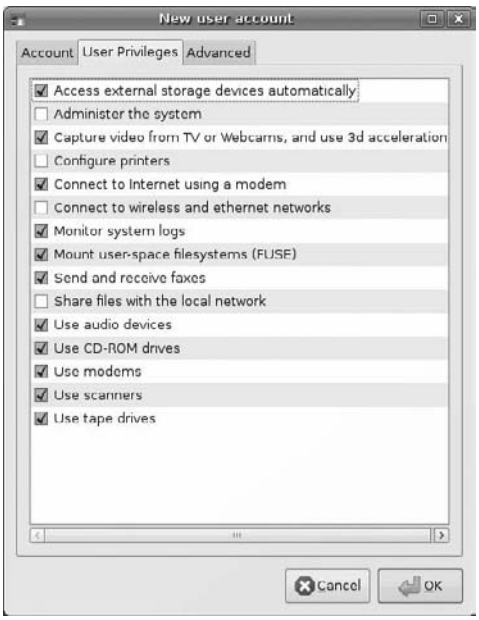


Only the Username and Password fields are mandatory on this dialog, but you should generally provide as much information as possible when creating a user account—you may not be the only administrator of the system that you're working on, so you should be kind to your fellow sysadmins who may one day wonder, "Who is this person?"

Aside from providing basic user information, one of the most interesting capabilities of this dialog is the ability to associate a specific profile with that user. A profile (not the same thing as the .profile file discussed in the section of Chapter 14 entitled "Exploring the Bash Configuration File") is a complete group of environment and privilege settings that is used when a new user is created. The standard profiles provided by the Users and Groups tool are named Administrator, Desktop user, and Unprivileged. Each of these profiles specifies things such as the devices that a user has access to, basic user capabilities, and the default groups that the user belongs to. You can determine the capabilities associated with any of the default user profiles by clicking the User Privileges tab, which displays the dialog shown in Figure B4-3.

FIGURE B4-3

The User Privileges tab



After checking the default privileges that are specified on this tab and making any changes that you want to make, click the Advanced tab to display the dialog shown in Figure B4-4.

Once you have finished customizing the account settings for the user account that you are creating, click OK on the main User Account Editor dialog to create the new account. The New User Account dialog closes, all system files (such as `/etc/passwd`, `/etc/group`, and `/etc/shadow`) are updated, and the new account is added to the list in the main Users Settings dialog.

FIGURE B4-4

The Advanced tab of the New User Account dialog



Managing Existing Users

In addition to simplifying the creation of new accounts, the Users and Groups tool also simplifies updating (or deleting) an existing account. To update or delete an existing account, start the Users and Groups tool by selecting the System ➪ Administration ➪ Users and Groups menu entry to execute this tool, and click Unlock to enable modification of your system user and group information. After you supply your password, the Add User and Delete buttons in the dialog shown in Figure B4-1 become active.

NOTE **The most common administrative task done after an account has been created is to add that account to other groups or remove it from groups of which it is currently a member. This is not done by modifying the settings for the account using the dialogs described in this section, but by using the procedures described later in this chapter in the section entitled "Managing Existing Groups."**

To modify or delete an existing account, scroll through the list of available accounts and select the account that you want to update or delete. Selecting the name of an existing account activates the Properties and Delete buttons.

To delete the selected account, click Delete. The Users Settings tool displays a dialog that requests confirmation that you actually want to delete the account, and that also explains that deleting an account just removes the entries for that user from the password, group, and shadow files—it does not remove any files and directories that are associated with the user, such as the user's home directory. To delete the

selected user, click Delete in the dialog. To return to the Users Settings tool without deleting the account, click Cancel. This also deselects the account that you had selected in the Users Settings tool.

To modify an existing account, click Properties. This displays essentially the same dialog and tabs as those in the New User Account dialog in the previous section, with a few differences:

- All dialogs are now titled "Account *user name* Settings" rather than "New User Account."
- Fields that cannot be edited, such as the user name and user ID, are disabled.
- You cannot change the profile associated with this account, because the profile is only used to provide default settings during account creation.

Once you have made any changes that you wish to make to the account, click OK to close the Settings dialog and return to the main Users Settings dialog. If you do not want to save any changes that you have made, click Cancel to return to the main Users Settings dialog.

Creating New Groups

Creating a new system group is usually done as a part of the installation process for any software that requires a group that does not yet exist on your Ubuntu system. System groups is the term used to refer to groups that are created to support a specific software subsystem or application, such as groups that control access to printers, databases, and various types of hardware. On Ubuntu and Kubuntu Linux systems, the group IDs of system groups are less than 1000.

You may want to create your own groups to do things like identify project members, to limit access to selected directories to a group, or to use a group to limit the people who can execute selected binaries. Ubuntu's Users and Groups tool provides a graphical interface to simplify both group creation and group management (the latter is discussed in the section after this one). To start this tool, select the System ➪ Administration ➪ Users and Groups menu entry to execute this tool. Click Unlock to display a dialog that enables you to supply your password (or that of another user with administrative privileges). After a password is supplied, the dialog shown in Figure B4-1 appears. To access the group-oriented portions of this tool, click the Manage Groups button, which displays a dialog like the one shown in Figure B4-5.

FIGURE B4-5

The Groups dialog in the Users and Groups tool



TIP **If the Groups Settings dialog shown in Figure B4-5 is completely grayed-out, you may have forgotten to click Unlock to enter an administrative password or mistyped that password. If this is the case, close the Groups Settings dialog, return to the Users Settings dialog, and click Unlock to redisplay the Authentication dialog. After entering a valid password for the administrative user that is selected on this dialog, click Authenticate to obtain the privileges necessary to edit users and groups. The Authenticate dialog closes automatically after a valid administrative password has been accepted.**

To create a new group, click Add Group. The dialog shown in Figure B4-6 appears.

FIGURE B4-6

The New Group dialog



Enter the name of the new group that you want to create in the Group name box. Next, you can optionally modify the group number associated with this group.

At this point, you can also add selected users to the group that you are creating by clicking the checkbox to the left of their names in the Group Members list. If you change your mind about adding a specific user to your new group, simply deselect the checkbox beside their name in the Group Members list.

Once you are finished changing settings and adding members to your new group, click OK to create that group and return to the Groups Settings dialog of the Users and Groups application. If you are finished creating or modifying groups, you can click Close to exit the Groups Settings dialog and return to the Users Settings dialog. To completely exit the Users and Groups application, click Close on the Users Settings dialog.

Managing Existing Groups

In addition to simplifying the creation of new groups, the Users and Groups tool also simplifies adding and removing users from those groups. To update groups' memberships, start the Users and Groups tool by selecting the System ➪ Administration ➪ Users and Groups menu entry to execute this tool. Click Unlock to display a dialog that enables you to supply your password (or that of another user with administrative privileges). After supplying your password, the dialog shown in Figure B4-1 appears. To access the group-oriented portions of this tool, click the Manage Groups button to display the dialog shown in Figure B4-5.

In this dialog, click the name of the group that you want to add or remove users from and click Properties. This displays a dialog that is essentially identical to the dialog shown in Figure B4-6 except that you cannot change the name of the group, only its members or Group ID.

TIP

You should rarely, if ever, use the group settings dialog to change group IDs. Changing the group ID in this dialog changes the entry for this group in the file `/etc/group` but does not change the group ID of any files that may happen to already be owned by that group. This could prevent any per-group protections that are already in use from working correctly, preventing users that should be authorized from accessing files or directories associated with their group.

To add a user to an existing group, click the checkbox to the left of that user's user name in the scrollable list at the bottom of this dialog. Similarly, to remove a user from an existing group, deselect the checkbox beside that user's user name in the scrollable list of group members.

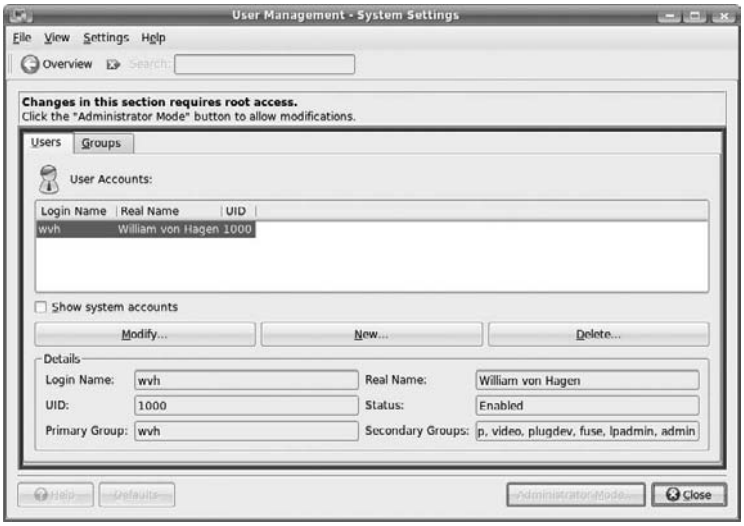
Any membership, name or group ID changes that you make in this dialog are not written to the /etc/group file (and updated in the `/etc/passwd` file, if necessary) until you click OK. You can click Cancel at any time to discard your changes. Clicking either of these buttons returns you to the Groups Settings dialog. Once you have finished making any changes to group memberships or other properties, click OK to close the Groups Settings dialog and return to the Users Settings dialog. To completely exit the Users and Groups tool, click Close.

Managing Users and Groups on KDE 3 Systems

Kubuntu 8.04 and other KDE 3–based systems provide a convenient graphical applet for creating and modifying users and groups as part of the System Settings tool. Select the K Menu ➪ System Settings menu command to start the System Settings application, and select the User Management configuration applet to execute this tool. The dialog shown in Figure B4-7 appears.

FIGURE B4-7

The KDE 3 User Management applet



In order to do anything other than simply see the lists of users and groups in this tool and view their properties, you must click Administrator Mode, enter your password in the dialog that appears, and click OK to proceed. This surrounds the active portion of this dialog with a red frame to indicate that you are working in a privileged mode, and enables the Modify, New, and Delete buttons on the Users and Group tabs.

Creating New Users

To create a new user in the User Management applet, click New to display the dialog shown in Figure B4-8.

FIGURE B4-8

Creating a new user on a KDE 3–based system



This dialog initially displays new_user in the Login name field as the default name of the user that you are creating, which you should replace with the actual login of the user that you want to create. Entering the actual user name will also update the name of that user's home directory to match the new login.

After entering the user's name and full name, click the Privileges and Groups tab to identify the groups to which the new user should belong. Figure B4-9 shows the Privileges and Groups tab.

As shown in Figure B4-9, new users have the privileges to access all of the storage and removable media devices attached to your system by default. If you do not want them to have these privileges, deselect the checkbox to the right of any permissions that you do not want them to have. They do not have administrative privileges by default (i.e., the ability to use the `sudo` command to perform administrative operations), which you can enable by clicking the Administer the system checkbox.

Next, click the Password & Security tab and enter an initial password for the new user in the Password field. Click OK to create the new user and return to the dialog shown in Figure B4-7. The new user will be displayed in the Users list in the center of this dialog.

FIGURE B4-9

Defining a new user's group memberships on a KDE 3–based system



Managing Existing Users

In addition to simplifying the creation of new accounts, the User management applet also simplifies updating (or deleting) an existing account. To update or delete an existing account if you are not already running this applet, select the K Menu ➪ System Settings menu command to start the System Settings application, and select the User Management configuration applet to execute this tool. The dialog shown in Figure B4-7 appears.

In order to do anything other than simply see the lists of users and groups in this tool and view their properties, you must click Administrator Mode, enter your password in the dialog that appears, and click OK to proceed. This surrounds the active portion of this dialog with a red frame to indicate that you are working in a privileged mode, and enables the Modify, New, and Delete buttons on the Users and Group tabs.

NOTE **The most common administrative task done after an account has been created is to add that account to other groups or remove it from groups of which it is currently a member. You can either do this by modifying the settings for the account using the dialogs described in this section or by using the procedures described later in this chapter in the section entitled "Managing Existing Groups."**

To modify or delete an existing account, scroll through the list of available accounts and select the account that you want to update or delete. To delete the selected account, click Delete. The User Management applet displays a dialog that requests confirmation that you actually want to delete the account, and provides a checkbox that enables you to simultaneously delete the primary group for that user, which generally has the same name as the user. To delete the selected user (and, optionally, the default group), click Delete in the dialog. To return to the Users Management applet without deleting the account, click Cancel.

To modify an existing account, click Modify. This displays the same dialog and tabs as those used when creating a new account, as described in the previous section, except for the fact that the user ID field is disabled because it cannot be edited at this point. To modify any aspect of this account, select the appropriate tab and make your changes.

Once you have made any changes that you wish to make to the account, click OK to close the Settings dialog and return to the main User Management dialog. If you do not want to save any changes that you have made, click Cancel to return to the main User Management dialog without committing your changes.

Creating New Groups

Creating a new system group is usually done as a part of the installation process for any software that requires a group that does not yet exist on your Kubuntu system. System groups is the term used to refer to groups that are created to support a specific software subsystem or application, such as groups that control access to printers, databases, and various types of hardware. On Ubuntu and Kubuntu Linux systems, the group IDs of system groups are less than 1000.

You may want to create your own groups to do things like identify project members, to limit access to selected directories to a group, or to use a group to limit the people who can execute selected binaries. The User Management tool provided on KDE 3–based systems such as Kubuntu 8.04 and earlier also provides a graphical interface to simplify both group creation and group management (the latter is discussed in the section after this one). If you are not already running this tool, select the K Menu ➪ System Settings menu command to start the System Settings application, and select the User Management configuration applet to execute this tool. The dialog shown in Figure B4-7 appears.

In order to do anything other than simply see the lists of users and groups in this tool and view their properties, you must click Administrator Mode, enter your password in the dialog that appears, and click OK to proceed. This surrounds the active portion of this dialog with a red frame to indicate that you are working in a privileged mode, and enables the Modify, New, and Delete buttons on the Users and Group tabs. To access the group-oriented portions of this applet, click the Groups tab, which displays a dialog like the one shown in Figure B4-10.

To create a new group, click New. The dialog shown in Figure B4-11 appears, using the first available group ID and displaying the default name new_group for the new group.

FIGURE B4-10

The User Management applet's Groups tab

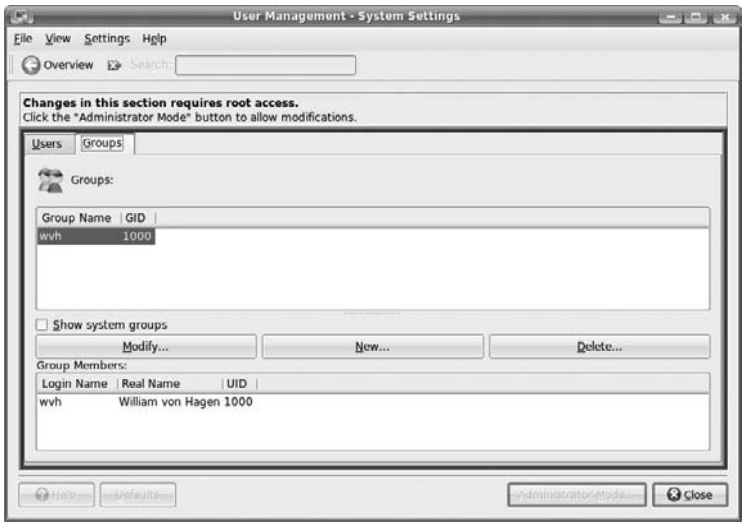


FIGURE B4-11

The New Group dialog



Enter the name of the new group that you want to create in the Group Name box. You can also optionally modify the group ID associated with this group.

At this point, you can also add selected users to the group that you are creating by selecting their names in the Available Accounts list and clicking Add to add them to the Selected Accounts list. If you change your mind about adding a specific user to your new group, select the name of that user in the Selected Accounts list and click Remove.

Once you are finished changing settings and adding members to your new group, click OK to create that group with the specified users and return to the Groups tab of the User management application. If you are finished creating or modifying groups, you can click Close to terminate the User Management applet and return to the System Settings application.

Managing Existing Groups

In addition to simplifying the creation of new groups, the User Management applet also simplifies adding and removing users from those groups. To update group memberships if you are not already running this applet, select the K Menu ➪ System Settings menu command to start the System Settings application, and select the User Management configuration applet to execute this tool. The dialog shown in Figure B4-7 appears.

In order to do anything other than simply see the lists of users and groups in this tool and view their properties, you must click Administrator Mode, enter your password in the dialog that appears, and click OK to proceed. This surrounds the active portion of this dialog with a red frame to indicate that you are working in a privileged mode, and enables the Modify, New, and Delete buttons on the Users and Group tabs. To access the group-oriented portions of this applet, click the Groups tab, which displays a dialog like the one shown in Figure B4-10.

In this dialog, click the name of the group that you want to add or remove users from and click Modify. This displays a dialog that is identical to the dialog shown in Figure B4-11. You can add selected users to the group that you are editing by selecting their names in the Available Accounts list and clicking Add to add them to the Selected Accounts list. You can remove users from the group that you are editing by selecting the name of that user in the Selected Accounts list and clicking Remove.

Once you are finished modifying the group that you are editing, click OK to save your changes and return to the Groups tab of the User management application. If you are finished creating or modifying groups, you can click Close to terminate the User Management applet and return to the System Settings application.

Managing Users and Groups on KDE 4 Systems

KDE 4, which is used on Kubuntu 8.10 and later systems, introduces many changes to traditional KDE applications and, in some cases, changes to the applications that are used for certain tasks. This is the case with the graphical application used to create and manage users and groups on KDE 4–based systems. As described in the previous section, these tasks used to be handled by an applet in the System Settings tool; they are now performed using a standalone application known as KUser.

Creating New Users

To create new users on a KDE 4–based system, select the K Menu ➪ Applications ➪ System ➪ User Manager menu command. A dialog appears, prompting you for your password. After entering your password and clicking OK, the User Manager application appears, as shown in Figure B4-12.

To create a new user, click the Add that appears at the far left in the KUser toolbar. A dialog appears that prompts you for the login name of the new user (not their full name). Enter the new user's login name and click OK to display the User Properties dialog shown in Figure B4-13.

FIGURE B4-12

The KDE 4 User Manager application, KUser



You use this dialog to provide any other information that you want to specify for the new user, such as their full name, office location and address, and so on. If you want the user to be able to log in using this account (in other words, if you are not creating a system account for some reason), you must deselect the Account disabled checkbox.

To specify the initial password for the new user that you are creating, click Set Password and enter the user's initial password in the dialog that appears—you must enter it twice to ensure that you have typed it correctly. Click OK to close the Enter Password dialog.

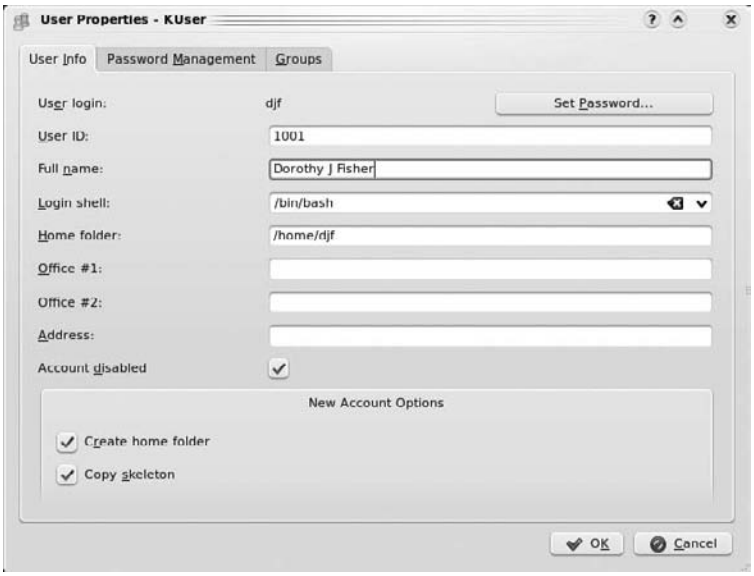
To specify the groups to which the user that you are creating belongs, click the Groups tab. At the time that this book was written, this dialog did not provide any default group memberships with the exception of the user's default group, which has the same name as the user. Some suggested group memberships are the following:

- To enable access to various storage devices and removable media, add the user to the **cdrom**, **dialout**, **fuse**, and **plugdev** groups.
- To enable the user to print, add them to the **lpadmin** group.
- To enable the user to see various log files in the `/var/log` directory, add them to the **adm** group.

- To enable the user to use the `sudo` command to obtain root privileges, add them to the **admin** group.
- If you have installed Samba on your system and want the user to be able to view shared directories in Nautilus and share their files via Samba, add them to the **sambashare** group. All users with admin group membership should also belong to this group.

FIGURE B4-13

Creating a new user in KUser



Before actually creating the new user, you can optionally use the Password management tab to configure the new user to conform to any password policies that you may use on your system(s), such as password expiration, mandatory password changes after a specified period of time, and so on. These are rarely necessary on personal or home systems, but you may be more paranoid than I am.

After completing the configuration of the new user account, click OK to create the account. You can then continue creating or modifying users and groups on your system, or select the File ➪ Quit menu command to exist the KUser application.

Managing Existing Users

Modifying existing users on a KDE 4–based system is done using the same graphical application that is used to create them. If you are not already running this application, select the K Menu ➪ Applications ➪ System ➪ User Manager menu command. A dialog appears, prompting you for your password. After entering your password and clicking OK, the User Manager application appears, as shown in Figure B4-12.

To edit an existing user, locate the user in the scrollable user list and click the entry for that user. The same dialog that is shown in Figure B4-13 appears; it contains the current information for that user. After making any modifications to the values on the User Info, Password Management, or Groups tabs, click OK to save the changes to that user's account, or click Cancel to close the User properties dialog without committing your changes.

Creating New Groups

Creating new groups on a KDE 4–based system is done by using the same graphical application that is used to create and modify users. If you are not already running this application, select the K Menu ➪ A pplications ➪ System ➪ User Manager menu command. A dialog appears, prompting you for your password. After you enter your password and click OK, the User Manager application appears, as shown in Figure B4-12. Click the Groups tab to display the list of existing groups, as shown in Figure B4-14.

FIGURE B4-14

The Groups tab in the KUser application



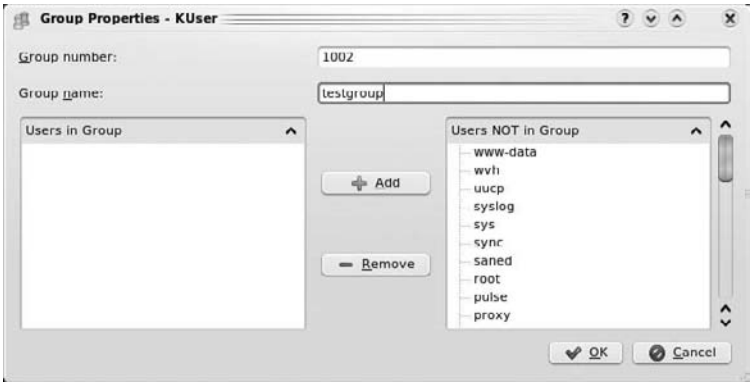
To add a new group, click the Add icon that appears in the Groups portion of the KUser toolbar (close to the center of the dialog). This displays the Group Properties dialog shown in Figure B4-15.

By default, this dialog displays the first free group ID that is not a system group (i.e., that is not less than 1000). You can change the group number by entering a new number in the Group number field. Next, enter the name of the group that you want to create in the Group name field.

To add users to the group that you are creating, select their logins from the scrollable User NOT in Group list, and click Add. If you change your mind, you can remove a user from the group by selecting their login in the scrollable Users in Group list and clicking Remove.

FIGURE B4-15

The Group Properties dialog in the KUser application



Once you have finished configuring your new group and its members, click OK to commit your changes and create the new group. If you decide that you do not want to create the new group, click Cancel. After selecting either of these buttons, the Group Properties dialog closes and the KUser application's main dialog reappears. If you are finished creating or configuring users and groups, select the File ➪ Quit menu command to close the KUser application.

Managing Existing Groups

Modifying existing groups on a KDE 4–based system is done using the same graphical application that is used to create them. If you are not already running this application, select the K Menu ➪ Applications ➪ System ➪ User Manager menu command. A dialog appears, prompting you for your password. After entering your password and clicking OK, the User Manager application appears, as shown in Figure B4-12. Select the Groups tab to display the list of groups that are available on your system, as shown in Figure B4-14.

To edit an existing group, locate the group in the scrollable group list and click the entry for that group. The same dialog shown in Figure B4-15 appears, containing the current information for that group and the list of its current members.

To add users to this group, select their logins from the scrollable User NOT in Group list, and click Add. To remove a user from this group, select their login in the scrollable Users in Group list, and click Remove.

Once you have finished modifying the group, click OK to commit your changes and create the new group. If you decide that you do not want to save your modifications, click Cancel. After selecting either of these buttons, the Group Properties dialog closes and the KUser application's main dialog reappears. If you are finished creating or configuring users and groups, select the File ➪ Quit menu command to close the KUser application.

PAMs and the Linux Authentication Process

Many different Linux and UNIX applications require authentication or special privileges of one type or another to access special devices or files, or to start processes as a specific user and group. In the early days of Linux/UNIX, each authentication-aware application was compiled with hardwired information about the authentication mechanism that it required. Changing or enhancing a system's authentication mechanism, therefore, required that all such applications be updated and recompiled, which is tedious at best even if you have the source code for all of the relevant applications on your system.

To resolve this problem, the folks at Sun Microsystems developed the idea of Pluggable Authentication Modules (PAMs), which provide a flexible and dynamic mechanism for authenticating any application or service that uses them. This model has since been adopted by most Linux and UNIX-like operating systems, including Ubuntu and Kubuntu Linux. Applications or services compiled with the Linux-PAM library use text-format configuration files to identify and describe their authentication requirements and the specific shared library modules used to implement them. The PAM model lets you modify the authentication requirements of existing applications by simply adding entries to the PAM configuration file that is used by a specific application or service. Applications that use PAMs are typically referred to as PAM-aware applications, and PAMs are usually simply referred to as PAM modules.

PAMs satisfy different parts of the authentication requirements for PAM-aware applications, much like reusable code and libraries do for applications in general. For example, a PAM-aware version of the login program can invoke several PAMs that check things such as whether the user logging in as root is on a terminal listed as a secure terminal, whether users are allowed to log in on the system at the moment, and other similar, authentication requirements. Because PAMs are shared library modules, a PAM-aware version of the ssh command can reuse the same "are users allowed to log in on the system now" PAM as the PAM-aware version of login, but then apply other rules that are more relevant to `ssh` than to login.

The files that describe the sequence of PAMs associated with specific applications or services are located in the directory `/etc/pam.d`. The shared-library PAM modules themselves are stored in the directory `/lib/security`. In older PAM implementations, the PAM modules used by all applications on a system were defined in a central configuration file, `/etc/pam.conf`. This file is still used, but only as a fallback if no PAM configuration file for a specific application is found in `/etc/pam.d`.

Many different PAM libraries are available in the Ubuntu repositories, some of which are installed by default and many that are associated with specific authentication mechanisms such as Lightweight Directory Access Protocol (LDAP) or Kerberos. Searching the repository for packages that match the string "libpam" and extracting the names of unique package names shows that there are (currently) 48 PAM-related packages on a Ubuntu or Kubuntu 8.10 system, as the following command shows:

```
$ apt-file search libpam | sed -e 's;:.*;;' | grep libpam | \
  grep -v dev | sort | uniq |  wc -l
48
```

NOTE **The `apt-file` command is not installed by default on Ubuntu systems, but can easily be installed using your favorite package management software. Installing the `apt-file` command was recommended in the discussion of useful package management utilities in Chapter 23, "Adding, Removing, and Updating Software."**

The PAM modules used by specific, nonstandard authentication mechanisms are identified as requirements for the packages that provide the applications and daemons associated with those authentication mechanisms and are, therefore, only installed when those authentication mechanisms are installed.

PAM Configuration Files for Applications and Services

The PAM configuration files in `/etc/pam.d` have the same name as the PAM-aware application or service that they are associated with and define the sequence of PAM modules invoked, in order, during the authentication and validation process for that application or service. In addition to PAM rules and related statements, the files in `/etc/pam.d` can contain blank lines (which are ignored) and comments—any characters on a line that follow the traditional hash mark (#) are interpreted as a comment.

Each noncomment line in one of the files in `/etc/pam.d` either is an `@include` statement that includes (inserts) the contents of another PAM definition file at this point or is an actual PAM rule for how a specific PAM module is used as part of the authentication process for the associated application or service. Included files are used to simplify PAM configuration files, enabling a common set of PAM rules to be included by other PAM configuration files. Each PAM rule can consist of four fields separated by white space, the first three of which are mandatory. These fields have the following meaning and content:

- `module-type:` The type of PAM module defined on that line. A module's type defines how a specific PAM module is used during the authentication process. Valid values are the following:
 - `auth:` Identifies modules that verify user identity or that system requirements have been met. Common system requirements are that a service can be started at the current time (for example, that `/etc/nologin` does not exist when a user is trying to log in), that an acceptable device is being used (the device is listed in the `file /etc/securetty`), that the user is already the root user, and so on.
 - `account:` Identifies modules that verify that the user can be authenticated based on system requirements such as having a valid account that is able to log in at the current time (based on per-account policies, not general system constraints), has access to the requested application or service, and so on.
 - `password:` Identifies modules that verify that a user can authenticate to the system.
 - `session:` Identifies modules associated with tasks that must be done before the associated service or application is activated, or just before the termination of that service or application. Modules of this type typically perform system functions such as mounting directories, logging audit trail information, or guaranteeing that system resources are available.
- `control-flag:` The control-flag specifies the implications of different return values from the specified PAM module. Valid values are the following:
 - `optional:` Indicates that success of the PAM module is not critical to the application or service unless this PAM is the only PAM for a specified module type. If it is, its success or failure determines the success or failure of the specified module type.
 - `required:` Indicates that success of the PAM module is mandatory for the specified module type. The failure of any PAM marked as required for a specific module type (such as all modules labeled as auth) is reported only after all required PAMs for that module type have been executed.
 - `requisite:` Indicates that failure of the PAM module immediately returns failure to the associated application or service.
 - `sufficient:` Indicates that success of the PAM module satisfies the authentication requirements of this module type. If no previous required PAM has failed, no other PAMs for the associated module type are executed after a module labeled as sufficient has succeeded. Failure of a PAM identified as sufficient is ignored as long as any subsequent modules that are identified as

required for that module type return success. If a previous required PAM has failed, the success of a PAM marked as sufficient is ignored.

- `module-name:` The name of the PAM module associated with this entry. By default, PAM modules are located in `/lib/security`, but this field can also identify modules located in other directories by specifying the absolute path and filename of a PAM module.
- `arguments:` Optional, module-specific arguments.

The next section provides an example of the structure and organization of a PAM configuration file by examining the PAM configuration file used by the login process on Ubuntu systems.

Example: PAMs Used by the Login Process

To see how a PAM file actually works, let's look at an actual example. The configuration file for the PAMs used by the login program is the `file /etc/pam.d/login`. After removing blank lines and comments, the meaningful parts of this file on an Ubuntu system are the following:

```
1:  auth      requisite  pam_securetty.so
2:  auth      requisite  pam_nologin.so
3:  session   required   pam_selinux.so close
4:  session   required   pam_env.so readenv=1
5:  session   required   pam_env.so readenv=1 envfile=/etc/default/
locale
6:  @include  common-auth
7:  auth      optional   pam_group.so
8:  session   required   pam_limits.so
9:  session   optional   pam_lastlog.so
10: session   optional   pam_motd.so
11: session   optional   pam_mail.so standard
12: @include  common-account
13: @include  common-session
14: @include  common-password
15: session   required   pam_selinux.so open
```

NOTE **The actual `/etc/pam.d/login` file does not include line numbers (and cannot). I've added the line numbers to make it easier to match up the PAM entries with the explanations that follow.**

The respective lines of this sequence of PAMs do the following:

1. Invokes the PAM module `pam_securetty.so` to check whether the user is logged in on a secure terminal as defined in the `file /etc/securetty`. This check must succeed, but is not sufficient to enable a user to log in.
2. Invokes the `pam_nologin.so` PAM module to check whether logins are allowed on the system at the current time, which is usually done by checking for the `file /etc/nologin`. If this file is present, the system is in a maintenance state or in the process of being shut down, and logins are disabled. This check must succeed (i.e., the file must not exist), but it is not sufficient to enable a user to log in.
3. Invokes the `pam_selinux.so` PAM module to terminate any SELinux security context that might be active. IF SELinux is not installed, this module simply returns success.

4. Invokes the `pam_env.so` module, which reads a list of environment variables to set or unset from the `file /etc/security/pam.conf` (which does not contain any noncommented entries on an Ubuntu system and is therefore somewhat meaningless). This PAM must succeed (i.e., be able to read its configuration file) but is not sufficient to enable a user to log in.
5. Invokes the `pam_env.so` module again, which in this case reads a list of environment variables to set or unset from the `file /etc/default/locale`, which sets the `LANG` environment variable on Ubuntu and Kubuntu systems. This PAM must succeed (i.e., be able to read its configuration file), but is not sufficient to enable a user to log in.
6. Includes the contents of the `file /etc/pam.d/common-auth`, which provides a common set of PAM rules for system authentication. On Ubuntu and Kubuntu systems that use standard `/etc/passwd` and `/etc/shadow` password authentication, this file contains the following PAM rules:

```
auth    [success=1 default=ignore]    pam_unix.so nullok_secure
auth    requisite                     pam_deny.so
auth    required                      pam_permit.so
```

 The first rule invokes the `pam_unix.so` to verify that the user has a valid entry in the `/etc/shadow` file. The `nullok_secure` argument indicates that null (empty) passwords are acceptable on this system when coming from the terminals listed in `/etc/securetty`. This PAM must succeed but is not sufficient to enable a user to log in. If this rule succeeds, the other rules in this file are not invoked. If the first rule fails, the second rule immediately returns failure to prevent authentication. The third rule puts a positive return value on the auth stack if no value is already present there, but should not be reached.

 On Ubuntu systems where other authentication mechanisms, such as LDAP, Kerberos, or fingerprint swiping have been installed, this file would invoke the PAMs specific to checking authentication as required by those services.
7. Invokes the `pam_group.so` module, which verifies that the user is a member of one or more groups and that a group with the numeric identifier for the user's login group actually exists in the `/etc/group` file. This PAM module is executed but is optional—its success or failure does not matter unless it was the only auth module for the login service (which it is not), and it is therefore not sufficient to enable a user to log in.
8. Invokes the `pam_limits.so` module to set user limits specified in the configuration file `/etc/security/limits.conf` (which does not contain any noncommented entries on an Ubuntu system and is therefore somewhat meaningless). This PAM must succeed (i.e., be able to read its configuration file) but is not sufficient to enable a user to log in.
9. Invokes the `pam_lastlog.so` module to add an entry for the user to the file `/var/log/lastlog`, adding an entry when opening a session and terminating that entry when a session is closed. This PAM module is executed but is optional—its success or failure does not matter unless it was the only session module for the login service (which it is not), and it is therefore not sufficient to enable a user to log in.
10. Invokes the `pam_motd.so` module to display the contents of the `/etc/motd` file on the login console. This PAM module is executed but is optional—its success or failure does not matter unless it was the only session module for the login service (which it is not), and it is therefore not sufficient to enable a user to log in.
11. Invokes the `pam_mail.so` module to examine the user's mail directory and display a message if there is any pending mail. This PAM module is executed but is optional—its success or failure

does not matter unless it was the only session module for the login service (which it is not), and it is therefore not sufficient to enable a user to log in.

12. Includes the contents of the file `/etc/pam.d/common-account`, which provides a common set of PAM rules for verifying that the user has a valid account on the system. On Ubuntu and Kubuntu systems that use standard `/etc/passwd` and `/etc/shadow` password authentication, this file contains the following PAM rules:

```
account [success=1 default=ignore]   pam_unix.so
account requisite                    pam_deny.so
account required                     pam_permit.so
```

The first rule invokes the `pam_unix.so` module to perform the traditional Linux authentication mechanism check that the user's account is not marked as expired in the `file /etc/shadow`. This PAM sets the account stack value to 1 on success but is not sufficient to enable a user to log in. If this rule succeeds, the other rules in this file are not invoked. If the first rule fails, the second rule immediately returns failure to prevent authentication. The third rule puts a positive return value on the auth stack if no value is already present there, but should not be reached.

On Ubuntu systems where other authentication mechanisms, such as LDAP, Kerberos, or fingerprint swiping have been installed, this file would invoke the PAMs specific to verifying accounts as required by those services.

13. Includes the contents of the file `/etc/pam.d/common-session`, which provides a common set of PAM rules for interactive and noninteractive login sessions. On Ubuntu and Kubuntu systems that use standard `/etc/passwd` and `/etc/shadow` password authentication, this file contains the following PAM rules:

```
session   [default=1]    pam_permit.so
session   requisite      pam_deny.so
session   required       pam_permit.so
session   required       pam_unix.so
session   optional       pam_ck_connector.so nox11
```

The first rule invokes the `pam_permit.so` module to permit session access to the system and sets the session stack value to 1 if no value has previously been set. If the first rule fails, the second rule immediately returns failure to prevent session authentication. If this rule fails for some reason, the third rule puts a positive return value on the auth stack if no value is already present there, but should not be reached. The fourth rule invokes the `pam_unix.so` module to create a log entry to the syslog service identifying the beginning of a user's login session on the current system. This PAM must succeed but is not sufficient to enable a user to log in. The last rule invokes the `pam_ck_connector.so` module to register a non-X11 login session with the ConsoleKit daemon, meaning that the daemon recognizes that a login on the physical console has occurred.

14. Includes the contents of the file `/etc/pam.d/common-password`, which provides a common set of PAM rules for actually checking a password. On Ubuntu and Kubuntusystems that use standard `/etc/passwd` and `/etc/shadow` password authentication, this file contains the following PAM rules:

```
password  [success=1 default=ignore] pam_unix.so obscure sha512
```

```
password  requisite                     pam_deny.so
password  required                      pam_permit.so
```

This rule invokes the `pam_unix.so` module to check `/etc/passwd` for a valid user and test the password for that user against the encrypted password entry found in the file `/etc/shadow`. The obscure argument specifies that, when changing a password, certain checks will be performed, such as verifying that the new password is not a palindrome or rotated version of the old password, that the new password does not differ from the old one simply in terms of case, and that the new password is not "too similar" to the old one. (I personally have never looked at the code to see how similarity is determined, but its rules seem to be fairly obscure.) The `sha512` option specifies the encryption algorithm that will be used to encrypt the password. This PAM must succeed but is not sufficient by itself to enable a user to log in. If this rule succeeds, the other rules in this file are not invoked. If the first rule fails, the second rule immediately returns failure to prevent authentication. The third rule puts a positive return value on the password stack if no value is already present there, but should not be reached.

15. Invokes the `pam_selinux.so` module with the open parameter to open a new SELinux security context for the new login session.

NOTE **Some other password checking and password-strength checking PAMs are listed in the file `/etc/pam.d/common-password` but are commented out. You may want to uncomment these and comment out the original `pam_unix.so` entry if, for example, you want to take advantage of the dictionary checks performed by the `pam_cracklib.so` module, which are irritating but ensure that a password is not a word that is commonly found in a dictionary.**

Once all of the PAMs specified in the `/etc/pam.d/login` file (and files that it included) have completed successfully, the login process continues, the user's login shell is created, and any other startup actions specified in the user's login configuration files are executed.

The PAM authentication process can be complex and conducive to migraine headaches. Even though it is a pain, security is one of any sysadmin's most important responsibilities. You will rarely, if ever, want to muck with the PAM configuration files on a home Linux system, but you may want to modify them (and should certainly understand them) when using Ubuntu or Kubuntu in an enterprise or academic environment. If it's any consolation, think how complex the code to implement all of this would have been without the flexibility that PAMs provide.

Configuration Files for Various PAMs

The text-format files in `/etc/pam.d` control the PAMs associated with each authentication-aware application or service. As mentioned when exploring the `/etc/pam.d/login` file in the previous section, some of these PAMs use optional configuration files to further refine their behavior. The configuration files for individual PAMs are located in the directory `/etc/security`. Although these files must exist, they do not need to contain any useful information—they are there in case you want to take advantage of the advanced configuration options that they provide. Here is a list of the files in this directory which are found on a Ubuntu or Kubuntu 8.10 system:

- `access.conf`: Provides fine-grained access control for logins and is used by the `pam_access.so` module.
- `group.conf`: Provides per-session group membership control and is used by the `pam_group.so` module.

- `limits.conf`: Provides a per-user mechanism for setting system resource limits; used by the `pam_limits.so` module.
- `namespace.conf`: Defines the directories that will be polyinstantiated if the `pam_namespace.so` module is being used to support the SELinux Multi Level Security (MLS) policy. Polyinstantiated directories are typically system temporary directories and user home directories. Polyinstantiation and namespaces mean that different users who are viewing a polyinstantiated resource will each see different contents based on user name, role, security level, and so on.
- `namespace.init`: Contains a shell script that will be executed if the `pam_namespace.so` module is being used to support the SELinux Multi Level Security (MLS) policy. This default script provided on Ubuntu and Kubuntu systems guarantees that the contents of the `/etc/skel` directory will be copied to polyinstantiated home directories when a new user account is created.
- `opasswd`: Contains a list of hashes of recently used passwords for various users to prevent password reuse. This file really isn't a PAM configuration file per se, but is used by PAMs such as `pam_cracklib.so`. This file is only used if you add the remember=*N* option to the pam_unix.so configuration line in the `/etc/pam.d/common-password` file, where *N* is the number of previous password hashes that you want the system to remember for each user.
- `pam_env.conf`: Provides a mechanism for setting environment variables to specific values; used by the `pam_env.so` module.
- `time.conf`: Provides a mechanism for imposing general or user-specific time restrictions for system services and applications; used by the `pam_time.so module`.

If you install other authentication systems and their associated PAMs, other configuration files associated with those PAMs may also be installed in `/etc/security`. This section only lists the default configuration files provided by a standard Ubuntu installation.

What If PAM Configuration Files Are Missing?

PAMs provide a very powerful mechanism for authenticating various applications and making sure that only the right users are running them from the right devices. It is, therefore, easy to see that the correct configuration of application-specific PAMs is very important. But what if an application uses PAMs and you forget to create its PAM configuration file or that file is somehow accidentally deleted? To cover these cases, the Linux-PAM library provides a default configuration file for any applications and services that do not have their own configuration files. This is the file `/etc/pam.d/other`, which has the following valid entries on a standard Ubuntu and Kubuntu 8.10 system:

```
@include common-auth
@include common-account
@include common-password
@include common-session
```

The contents of each of these included files was discussed earlier in this chapter, in the section entitled "Example: PAMs Used by the Login Process." On Ubuntu and Kubuntu systems, by default, the `/etc/pam.d/other` PAM configuration file provides a reasonable set of defaults for verifying that random applications are at least being executed by someone who can successfully log in and authenticate him- or herself to the system, and guarantees that various log entries for that user are created.

TIP **Because a missing PAM configuration file generally indicates a misconfigured system or that someone has imported a PAM-aware binary without thinking things through, you can use the `/etc/pam.d/other` file to disallow access and impose extremely paranoid security, by changing it to contain the following entries:**

```
auth       required   pam_deny.so
auth       required   pam_warn.so
account    required   pam_deny.so
account    required   pam_warn.so
password   required   pam_deny.so
password   required   pam_warn.so
session    required   pam_deny.so
session    required   pam_warn.so
```

Because subsequent required entries for a given module type are still executed, each module-type entry first executes the `pam_deny.so` PAM, which denies access to the requested service, and then also executes the `pam_warn.so` PAM, which enters a warning message to the system log. These log entries will help you identify the problem if a cranky user hasn't already surfaced and asked why they can no longer run some obscure binary that they've imported.

Customizing the sudo Command

As introduced in the section of Chapter 4 entitled "Performing Privileged Operations," Ubuntu and Kubuntu do not use the traditional root account to perform privileged operations but instead enable users who are members of the admin group to perform all privileged operations by specifying their own passwords. The rationale for why the root account (and its traditional friend, the `su` command) are disabled on these systems is discussed in the official online Ubuntu documentation at `https://help.ubuntu.com/community/RootSudo`. To save you a Web lookup, this is basically viewed as a security improvement, which it certainly can be. Regardless of whether you like this approach, hate it, or are simply puzzled by it, that's the way that Ubuntu and Kubuntu Linux work. Even if you passionately believe that this is an odd approach, you are probably not going to be able to persuade the entire Ubuntu and Kubuntu community that it is wrong. Save your breath.

TIP **The page at `https://help.ubuntu.com/community/RootSudo` explains how to change a Ubuntu or Kubuntu system to reenable the traditional root account. Although this is possible, I strongly suggest that you do not do this unless you have an excellent, site-specific reason to do so and are willing to have your Ubuntu and Kubuntu system(s) be fundamentally different from all other Ubuntu and Kubuntu systems in the known universe.**

This section explains the internals of the `sudo` command and its configuration file, and discusses how to grant administrative privileges to other users in various forms.

The following list highlights some of the ways in which the `sudo` command works and should be used on Ubuntu and Kubuntu systems:

- `sudo` is a command that enables you to execute privileged command-line utilities, such as `apt-get`, `aptitude`, `dpkg`, and so on. If you want to execute graphical commands with administrative privileges on a standard Ubuntu system, you should use the `gksudo` command instead. If you want to execute graphical commands with administrative privileges on a standard Kubuntu system, you should use the `kdesudo` command instead.

- The problem with running graphical applications using sudo rather than `gksudo` or `kdesudo` is that they may change the ownership of the `.ICEauthority` file in your home directory, which many KDE, GNOME, and Standard X Window system applications use to track the applications that are authorized to access a specific X Window system display. If you accidentally use sudo to start a graphical application and cannot execute subsequent graphical apps, you can remove the `.ICEauthority` file in your home directory by using the command `sudo rm .ICEauthority`, and then use `gksudo` to start your graphical applications.
- Once you have executed the `sudo`, `gksudo`, or `kdesudo` commands to run any privileged commands, your authorization is stored for 15 minutes so that you don't have to provide your password over and over for each privileged command that you want to execute. During this time, you must still supply the `sudo`, `gksudo`, or `kdesu` commands, but you will not be prompted for a password. After 15 minutes, using `sudo`, `gksudo`, or `kdesudo` again will result in your being prompted for a password, and the cycle repeats.

TIP **If you are really lazy (as I have been known to be), a quick hack for using `sudo` is to use the `sudo` command to execute the bash shell by executing the command `sudo -s`. After supplying your password, this will give you a root shell in which you no longer have to specify the `sudo`, `gksudo`, or `kdesu` commands to execute privileged operations. This is convenient but is bad form and will generate verbal abuse from any other Ubuntu and Kubuntu fans that see you doing this. Please be clandestine if you should ever want to do this.**

The operation of the `sudo` command is controlled by the file `/etc/sudoers`, which is a text file that cannot be viewed by mere mortals for security reasons. To view this file, you should therefore use the `visudo` command to display it, using a command like the following:

```
$ sudo visudo
```

The `visudo` command is a special version of the `vi` command that ensures that only one person is modifying the `/etc/sudoers` file at a time. You, therefore, do not need to identify the file `/etc/sudoers` as an argument—the `visudo` command knows what file it is designed to edit. You can, of course, open this file in any text editor, but you won't get the protection against simultaneous edits that the `visudo` command provides.

As you will see, the `/etc/sudoers` file can contain blank lines (which are ignored) and comments—any characters on a line that follow the traditional hash mark (#) are interpreted as a comment. The meaningful entries in the `/etc/sudoers` file on a standard Ubuntu or Kubuntu system are the following:

```
Defaults   env_reset
root        ALL=(ALL) ALL
%admin      ALL=(ALL) ALL
```

The first of these lines specifies settings that identify default options for how the `sudo` command behaves. The only option in the Ubuntu or Kubuntu `/etc/sudoers` file is the following:

- `env_reset`: Resets the environment to only contain the `LOGNAME`, `SHELL`, `USER`, `USERNAME`, and `SUDO_*` environment variables. This prevents the propagation of other environment variables to privileged environments.

One option that you may want to change in the `/etc/sudoers` file is the default period of time for which `sudo`, `gksudo`, and `kdesudo` cache the credentials of the root user so that you don't have to enter your password again within a set period of time. By default, these commands cache credential

for 15 minutes, which many administrators view as being too long. To explicitly set this cache period (known as the timeout period), modify (or add) the Defaults entry in your `/etc/sudoers` file to look like the following:

```
Defaults    env_reset, timestampt_timeout=XX
```

Change *XX* to be the number of minutes for which you want your privileges to be cached without requiring you to enter your password again. To require all users to always enter a password when entering any of these commands, set *XX* to 0.

The next two lines are the really interesting ones. The following entry states that root can do anything and execute any command on the system:

```
root      ALL=(ALL) ALL
```

Fixing a Broken sudo Command

Installing some versions of Ubuntu in expert mode uses the standard root mechanism, which can leave your system in a state where normal users cannot use the `sudo` command to execute privileged commands. Similarly, even on a normal Ubuntu system, somehow losing the contents of your `/etc/group` or `/etc/sudoers` files can leave anyone unable to perform any privileged operation. In the latter case, you can reboot your system, press the Escape key to see the GRUB boot menus, and select the Recovery Mode entry from the GRUB menu. This will reboot your system and bring up a root shell.

Once you are logged in as root on your Ubuntu system, check to see that your `/etc/group` file contains the admin group and that any users who you want to be able to perform administrative operations are indeed members of that group. If you do not have an admin group, you can add it using the Users and Groups tool explained earlier in this chapter or by a command-line command (as the root user) such as the following:

```
# addgroup --system admin
```

Next, make sure that the `/etc/sudoers` file contains the entry that grants administrative privileges to the admin group, which is:

```
%admin ALL=(ALL) ALL
```

If this entry is not present in your `/etc/sudoers` command, you can use the `visudo` command to edit that file and add the entry.

After making sure that the admin group exists and is referenced in the `/etc/sudoers` file, all you have to do is add any user that you want to be able to run protected applications to the admin group. You can do this by using the Users and Groups tool discussed earlier in this chapter or from the command-line by using a command such as the following:

```
# adduser djf admin
```

This command would add the user `djf` to the admin group, enabling them to run privileged operations using `sudo` and `gksudo` (or `kdesu`) in the future.

The following line states that all members of the admin group (as defined in the `/etc/group` file) can also do anything (i.e., run any command as `root`):

```
%admin    ALL=(ALL) ALL
```

To enable other users to perform privileged operations, all you have to do is add them to the admin group in the `/etc/group` file using the Users and Groups utility (as explained earlier in this chapter) or by using the `sudo` command to execute a command-line utility such as `adduser`, as in the following example:

```
$ sudo adduser djf admin
```

This command would add the user `djf` to the admin group and enable that user to execute any commands using `sudo` or `gksudo`.

Using ACLs for Sophisticated File Sharing

Standard Linux/UNIX file permissions are fine if you have a relatively small number of users with limited requirements for sharing and working on the same files. The section of Chapter 4 entitled "Basic Concepts: Users and Groups" introduced the idea of groups, while the section of Chapter 4 entitled "File and Directory Permissions under Linux" introduced file and directory permissions and explained the portions permissions that are relevant to the group that owns a file or directory.

Linux groups are the traditional way of enabling access to selected files and directories by multiple users. Unfortunately, using groups to control shared access requires the intervention of a system administrator every time you want to add a new user to the group or remove an existing user, as well as set up the initial shared directories. Access Control Lists (ACLs), which are supported in most modern versions of Linux, including Ubuntu Linux, eliminate this hassle by providing a fine-grained set of permissions that users can impose on their own directories and that go far beyond the permissions and protections provided by standard Linux groups. ACLs have been around for a long time on Linux and other UNIX-like systems but have never become as popular as they might be, because of the arcane syntax of the command-line tools that are traditionally used to view and manage them. As you'll see later in this chapter, graphical tools for working with ACLs are now available, and Ubuntu and Kubuntu users can now set and modify ACLs with ease.

Simply put, an ACL is a list of Linux users and/or groups and the access rights that they have to a specific file or directory. ACLs enable you to define totally granular permissions such as "only the users wvh and djf can write this file, but the user juser can at least read it" without requiring that you create any special-purpose Linux groups. The key item of interest about ACLs is that you can set and manage the ACLs on any files that you own, giving various users and groups access to those files without requiring the intervention of anyone with administrative privileges. ACLs truly put you in the driver's seat as far as file and directory access goes.

If you're playing a round of acronym bingo, Access Control Lists (ACLs) as implemented on Linux systems today are defined by the draft Portable Operating System Interface (POSIX) standard 1003.1e, draft 17, from the Institute of Electrical and Electronics Engineers (IEEE). This is not an official standard, but it is publicly available and has become the foundation for ACL implementations for modern operating systems such as Linux.

The rest of this chapter explains how to install the commands necessary to use ACLs, other system requirements that may be associated with using ACLs on your systems, and how to use various command-line and graphical tools to examine, set, update, and remove ACLs in your files and directories.

NOTE **This section does not discuss extended file attributes, which are often presented as a companion topic to ACLs, although they do something completely different.**

Overview of Linux ACLs

Linux supports two basic types of ACLs:

- ACLs used to control access to specific files and directories
- Per-directory ACLs known as *mask* ACLs that define the default ACLs that are assigned to any files created within a directory

ACLs are represented in a standard format consisting of three colon-separated fields:

- The first field of an ACL entry is the entry type, which can be one of user (u), group (g), other (o), or mask (m).
- The second field of an ACL entry is a user name, numeric User ID (UID), group name, or numeric Group ID (GID), depending on the value of the first field. If this field is empty, the ACL refers to the user or group that owns the file or directory. This field is empty for ACLs of the mask and other types.
- The third field is the access permissions for this ACL. These are represented in two forms:
 - A standard UNIX-like permissions string rwx (Read, Write, and eXecute permissions, where eXecute permissions on directories indicate the ability to search that directory). Each letter may be replaced by a - (dash), indicating that no access of that type is permitted). These three permissions must appear in this order.
 - A relative symbolic form that is preceded by a + (plus) sign or a ^ (caret) symbol, much like the symbolic permissions that are designed for use with the `chmod` command by people who are octally challenged. In this ACL representation, the + or ^ symbols are followed by single r, w, or x permission characters, indicating that these permissions should be added to the current set for a file or directory (the plus symbol) or that these permissions should be removed from the current set (the caret) for a given file or directory.

When listed or stored in files, different ACL entries are separated by white space or new lines. Everything after a # character to the end of a line is a comment and is ignored.

Installing ACL Commands

Special-purpose commands such as those used to create and manage ACLs are installed as part of a default Ubuntu or Kubuntu installation. They are not installed by default on Ubuntu Server systems because they are less critical on systems to which few users have access. However, as with all software packages on Ubuntu, they can easily be installed on an Ubuntu server using your favorite package management software, as discussed in Chapter 23. The user-level commands for ACL support are located in the acl package.

If you expect people to actually use ACLs on your systems, you should also install the `eiciel` package, which is a graphical utility for setting and managing ACLs. Command-line ACL syntax can be tortuous, even for experienced system administrators—if you want people who are not command-line aficionados to actually use ACLs, the `eiceil` application will be a lifesaver. The `eiciel` application is discussed later in this chapter.

NOTE **Installing `eiciel` on a Kubuntu system will also install all of the dependencies required to run GNOME applications on your Kubuntu system, which will require around 425 MB of disk space. `eiciel` itself is fairly small, but its dependencies are large. I consider the increased chance that someone will actually use ACLs that is provided by `eiciel` to be well worth the disk space, but your mileage may vary.**

Activating Filesystem Support for ACLs

To use ACLs to enhance the granularity of permissions on your system, any filesystems that contain files and directories with which you want to use ACLs must be mounted with ACL support. This is not done by default on Ubuntu or Kubuntu systems but is easy enough to change. Filesystems mounted with support for ACLs will have the `acl` keyword in the mount options portion of your `/etc/fstab` file. Because this is not enabled by default, the entry in `/etc/fstab` for your filesystems will probably look something like the following:

```
/dev/hda1   /      ext3   relatime,errors=remount-ro      0 0
```

You can edit this file to enable ACL support when that filesystem is mounted by adding the `acl` keyword to the mount options for that filesystem. You can do this using your favorite text editor and a command like the following:

```
$ sudo vi /etc/fstab
```

Changing this entry to add support for ACLs would look like this:

```
/dev/hda1   /      ext3   relatime,acl,errors=remount-ro  0 0
```

Updating this file only means that support for ACLs will be enabled the next time your system automatically mounts that filesystem, which is typically the next time that you reboot your system. If you want to get started using ACLs right away, you can enable ACL support in the currently mounted version of the filesystem on `/dev/hda1` without rebooting your system, by executing a command such as the following, which would remount the example ext3 filesystem `/dev/hda1`, activating ACL support:

```
$ sudo mount -o remount /dev/hda1
```

Now, as they say in the Worldwide Wrestling Federation, "Get ready to ACL!"

Working with ACLs from the Command Line

As you'll see in this section, ACLs are ugly on the command line, often being indistinguishable from typing on a broken keyboard or displaying the contents of a binary file in a terminal window. Regardless of this, it is useful to understand their command-line representation and the command-line commands used to manipulate them. If you ever need to write a shell script to perform ACL manipulation, you'll be proud to say, "Of course I know how to use ACLs from a shell prompt."

Overview of Command-Line Utilities for ACLs

The Linux `acl` package provides the following three utilities for ACL creation, modification, and examination:

- `chacl`: Lets you change, examine, or remove user, group, mask, or other ACLs on files or directories.
- `getfacl`: Lets you examine file access control lists for files and directories.
- `setfacl`: Lets you set file and directory access control lists.

The next few sections explain how (and when) to use these commands.

Displaying Current ACLs from the Command Line

As an example of using ACLs, let's use a directory with the following contents and permissions:

```
$ ls -al
total 49
drwxr-xr-x    2 wvh wvh      80 2008-06-26 04:59 .
drwxr-xr-x  106 wvh wvh    5288 2008-06-26 23:11 ..
-rw-r-----    1 wvh wvh   44032 2008-06-26 04:58 resume.xml
```

The default ACL for this directory is the following:

```
$ getfacl .
# file: .
# owner: wvh
# group: wvh
user::rwx
group::r-x
other::r-x
```

The default ACL for the file `resume.xml` is the following:

```
$ getfacl resume.xml
# file: resume.xml
# owner: wvh
# group: wvh
user::rw-
group::r--
other::---
```

The default ACL for a file in a directory for which no default ACL has explicitly been set is derived from the default file protections associated with the user that created the file. A user's default file protections are derived from the user's umask, which is discussed in more detail in the sidebar in Chapter 14 entitled "Default Permissions When Creating Files and Directories."

Setting ACLs from the Command Line

There are three common ways to change the ACL of a file or directory:

- By explicitly setting it to a specific value using the `setfacl` command, which overwrites any existing ACL settings.
- By using the `setfacl` command with the `-m` (modify) option to modify an existing ACL.
- By using the `chacl` command to modify an existing ACL.

In the next few examples, I'll use the `chacl` command to change ACLs, because it is safer and easier to use because it doesn't overwrite any existing ACL. It also provides a bit more information about how ACLs really work than the `setfacl` command does.

For example, to add the user `djf` as someone who could read the file `resume.xml`, I would use a `chacl` (change ACL) command like the following:

```
$ chacl u::rw-,g::r--,o::---,u:djf:r--,m::rw- resume.xml
```

No, that isn't a typo or misprint—that's the way that ACLs look in real life. You wanted power, didn't you? It comes at a price. As mentioned previously, ACLs consist of three colon-separated fields that represent the permission of the user (owner of the file), group (the group ownership to the file), and others. When changing an ACL with the `chacl` command, you need to specify the ACL of the file, and then append the changes that you want to make to that ACL. The `u::rw-,g::r--,o::---` portion of the ACL in this example is the existing ACL of the file, and the `u:djf:r--,m::rw-` portion specifies the new user that I wanted to add to the ACL for that file and the effective rights mask to be used when adding that user's ACL. The effective rights mask is the union of all of the existing user, group, and other permissions for a file or directory. You must specify a mask when adding a new user to the ACL for a file.

Note that the name of the user that I am adding to this ACL appears between the two colons in the `u` portion of the ACL entry. Using the `getfacl` command to retrieve the ACL for my `resume.xml` file shows that the user `djf` has indeed been added to the list of people who have access to the file:

```
$ getfacl resume.xml
# file: resume.xml
# owner: wvh
# group: wvh
user::rw-
user:djf:r--
group::r--
mask::rw-
other::---
```

Using the `ls -al` command shows that the visible, standard UNIX file and directory permissions haven't changed:

```
$ ls -al
total 49
drwxr-xr-x    2 wvh wvh      80 2008-06-26 04:59 .
drwxr-xr-x  106 wvh wvh    5288 2008-06-26 23:11 ..
-rw-r-----    1 wvh wvh   44032 2008-06-26 04:58 resume.xml
```

Even more interesting and useful than just giving specific individuals read access to files is the ability to give specific users the ability to write specific files. For example, to add the user `djf` as someone who could both read and write the file `resume.xml`, I would use a `chacl` command like the following:

```
$ chacl u::rw-,g::r--,o::---,u:djf:rw-,m::rw- resume.xml
```

Note, again, that the name of the user that I am adding to this ACL appears between the two colons in the `u` portion of the ACL entry. Using the `getfacl` command shows that the user `djf` now has both read and write access to the file:

```
$ getfacl resume.xml
# file: resume.xml
# owner: wvh
# group: wvh
user::rw-
user:djf:rw-
group::rw-
mask::rw-
other::---
```

It's unfortunate that ACLs are so ugly, but that's because there is a fair amount of information packed into them. You may not find it odd that typos are common when typing ACLs. Now if there were only a friendly, graphical tool for ACL control . . .

Working with ACLs Using Graphical Tools

Yes, Virginia, there is a Santa Claus, and in this case his name is Roger Ferrer Ibanez. Roger is the primary author of the `eiciel` GNOME File ACL editor. What does *eiciel* mean? Try pronouncing it slowly with a fake French accent, and you'll hear that it sounds like Maurice Chevalier discussing his favorite computer science topic, the ACL.

The `eiciel` ACL editor provides a graphical interface for setting, updating, and removing ACLs in ext2, ext3, XFS, JFS, and ReiserFS filesystems, and should work equally well with any other filesystems that support POSIX ACLs.

You installed `eiciel` as part of the instructions in the section entitled "Installing ACL Commands on Ubuntu." Installing `eiciel` on a Ubuntu system creates a menu item for it on the Applications ➪ System Tools menu. Installing `eiciel` on a Kubuntu system creates a menu item for it on the Applications ➪ System menu.

If you are starting `eiciel` from the command line, you can supply the name of the file or directory for the ACLs you want to set or modify as a command-line argument. To open a file or directory from within `eiciel` itself, click Open in `eiciel's` upper-right corner, navigate to the file or directory of the ACLs you want to set or modify, and click Open to open the ACLs for that file or directory. Figure B4-16 shows the ACL for the file `resume.xml` (which was used as an example in the command-line discussion of ACLs earlier in this chapter).

FIGURE B4-16

The ACLs for a sample file in eiciel



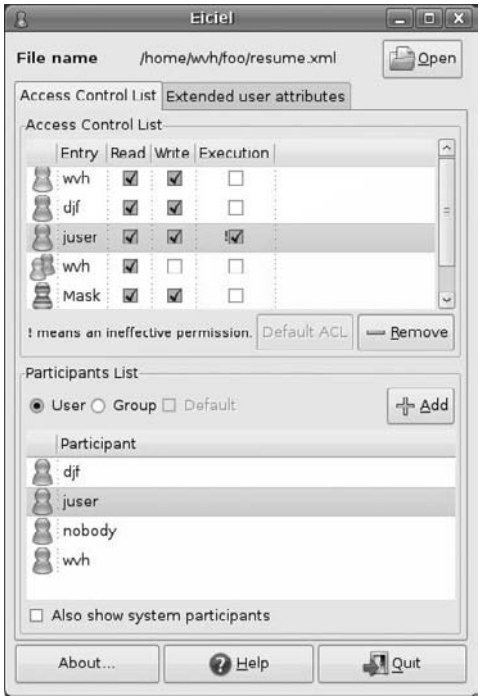
The top portion of the Eiciel dialog displays the current ACLs for the file or directory that you opened in `eiciel`. By default, the bottom half of the Eiciel dialog provides a list of all of the nonsystem users on your system. If you want to add a group to an ACL rather than a single user, select the Group radio button. To add system users and groups (i.e., those with UIDs and GIDs less than 1000), select the Also show system participants checkbox near the bottom of the Eiciel dialog.

To add a specific user or group to the ACLs for the file or directory that is currently displayed in `eiciel`, select the name of that user or group from the list in the bottom half of the Eiciel dialog. This activates the Add button `eiciel`.

To add the user to the ACL, click Add. The name of that user is added to the list in the top of the Eiciel dialog and is immediately added to the ACL for the file or directory that you are editing. Figure B4-17 shows the Eiciel dialog after the user `juser` has been added to the ACL for the file `resume.xml`.

FIGURE B4-17

The ACL entry for a newly added user



If you look closely at the ACL that I just added, you'll note that the Execute bit is labeled as being an ineffective permission. This is because the file itself isn't executable.

At this point, you can select the various settings for any user on the ACL displayed in the window at the top half of the Eiciel dialog. Any changes that you make to the ACL take place immediately—there is no need to explicitly save them. When you are finished using `eiciel`, click Quit to exit.

The `eiciel` editor is an excellent piece of software, and a great example of how the latest generation of Linux tools can remove much of the wizardry, typos, and barriers to use that are often the companions of complex command-line tools and twisted syntax. For more information about using `eiciel`, click Help in the Eiciel dialog to see its online manual. For information about the latest and greatest `eiciel` developments, see its home page at `http://rofi.pinchito.com/eiciel`. Thanks for a great piece of software, Roger!

Configuring Kerberos Authentication

Kerberos is a distributed authentication and communication service that was originally developed at the Massachusetts Institute of Technology (MIT). Kerberos provides secure authentication and communication for client/server applications by using strong cryptography to enable clients to prove their identity to servers over the network.

Kerberos works by exchanging encrypted security information between clients (which can be users or machines), the Kerberos Authentication Server, and the resource you are trying to access. The information that is initially exchanged when attempting to prove one's identity is known as a ticket. The information used to encrypt tickets and subsequent communications is known as a key. Once the identity of a client is verified, that client is granted a Kerberos token that can be used to verify their identity to any Kerberos-aware service. For security reasons, Kerberos tokens are timestamped so that they automatically expire unless renewed by a user or service. The primary system for granting tickets (and which houses the master copy of the Kerberos database) is known as the Kerberos Key Distribution Center (KDC).

The timestamps contained within Kerberos tokens (and tickets) can only be verified if the time and data are synchronized across Kerberos clients and servers. Kerberos authentication will fail if client and server clocks become skewed by more than 5 minutes. You should always run an NTP (Network Time Protocol) daemons on all Kerberos servers and NTP clients on all Kerberos clients to guarantee that their clocks remain in sync.

Kerberos uses the term “realm” to differentiate between authentication and Internet domains. A Kerberos realm is a set of machines that rely on a specific Kerberos server for authentication and that, therefore, trust that server. In Kerberos configuration files, your realm is typically identified in uppercase characters in order to differentiate it from any similar DNS domain with which it is associated.

MIT's Kerberos implementation is only one of several. Several alternate Kerberos implementations have been created over the years, usually to get around United States export restrictions that have since been lifted. For example, Novell’s SUSE (Software- und System-Entwicklung) Linux systems use an alternate Kerberos client/server implementation known as Heimdal (`www.pdc.kth.se/heimdal`). This section focuses on vanilla Kerberos from MIT as provided in the Ubuntu repositories.

Installing Kerberos Client and Server Commands

Using Kerberos requires that you either build and install it yourself from scratch or that a few basic packages are already installed on your systems. If you need to build it yourself, you can download the latest version from MIT at `http://web.mit.edu/kerberos/www`. However, building Kerberos is rarely necessary on Ubuntu and Kubuntu systems, because the Ubuntu repositories already provide an up-to-date version of the Kerberos packages that you’ll need on the system that acts as your Kerberos server and

on any clients of that Kerberos server. To install Kerberos, you must install the following packages using your favorite package management tool:

- `krb5-admin-server`: You must install this package on all Kerberos server and slave server systems. This package provides the programs that must be installed on a Kerberos 5 server or server replica.
- `krb5-config`: This package provides a template configuration file for Kerberos and associated documentation, and is automatically installed for you when you install the `krb5-admin-server` or `krb5-user` packages.
- `krb5-kdc`: This package provides the Kerberos Key Distribution Center (KDC) binaries, related utilities, a template configuration file, and associated documentation, and is automatically installed for you when you install the `krb5-admin-server` package.
- `krb5-user`: You must install this package on all client systems. This package provides basic Kerberos programs (`kinit`, `klist`, `kdestroy`, `kpasswd`) as well as Kerberized versions of the `telnet` and `ftp` applications.
- `libpam-krb5`: You must install this package on all Kerberos client and server systems in order to install the Kerberos PAM so that you can integrate Kerberos into the PAM authentication process on those systems.
- `ntp`: As mentioned in the introduction to this section, you should install a Network Time Protocol (NTP) server on your Kerberos server so that its timekeeping is accurate. The tickets used to authenticate clients with the Kerberos service are timestamped, and it is therefore extremely important that the clocks on the Kerberos server and all clients be synchronized.
- `ntpdate`: As mentioned in the introduction to this section, you should install a Network Time Protocol (NTP) client on all of your Kerberos client systems so that their timekeeping is accurate. The `ntpdate` client is installed by default on all Ubuntu and Kubuntu desktop systems—I'm just listing it here to make sure that you understand that it must be available on your system and should synchronize with the NTP server on your Kerberos server.

Installing the `krb5-admin-server` package displays a warning message that explains that you must actually configure it for your site before you can use Kerberos authentication. No big surprise there—read on!

Configuring NTP on Servers and Clients

Before configuring and starting your Kerberos server, you should first make sure that your Kerberos server has the correct time. You should also configure each client of that Kerberos server to get its time information from the Kerberos server before starting any Kerberos client. You can configure the client to use the Kerberos server's time server when you install and configure the Kerberos software on the client system or beforehand—the instructions for client NTP configuration are provided in this section to help ensure that you don't forget this step.

Configuring an NTP Server

Installing the `ntp` package on the system that will be your Kerberos server automatically configures your server to consult the time server at `ntp.ubuntu.com`. In case this host goes down or is otherwise unavailable, you may want to configure your time server to be able to consult one or more other time servers from the pool of public NTP servers that is available on the Internet. For a list of public NTP

servers and information about using them, see http://support.ntp.org/bin/view/Servers/WebHome.

The configuration information for an NTP server is stored in the configuration file /etc/ntp.conf. Edit this file in your text editor of choice, search for the default ntp.ubuntu.com entry, and add lines like the following after that entry:

```
server north-america.pool.ntp.org
server pool.ntp.org
```

If you are not located in North America (as I am), you would change the first line to be the name of a time server that is more appropriate for your geographic location, as explained at http://support.ntp.org/bin/view/Servers/NTPPoolServers.

After entering entries for additional time servers, save the file, exit your text editor, and restart the NTP server using a command like the following:

```
/etc/init.d/ntp restart
```

This will restart your time server with the new configuration entries.

Configuring an NTP Client

As mentioned previously, the ntpdate NTP client software is installed by default on all Ubuntu and Kubuntu desktop systems. When installing Kerberos on a client system, you should make sure that the NTP client software on each client system gets its time synchronization information from your NTP server. This ensures that the timestamps on Kerberos tickets are synchronized throughout your site.

Unlike most system services, the ntpdate client is actually run whenever you bring up an Ethernet interface. The script /etc/network/ifup.d/ntpdate is executed whenever you bring up an Ethernet interface. This script runs the /usr/sbin/ntpdate-debian script, which in turn executes the /usr/sbin/ntpdate program with the timeservers that are specified in the /etc/default/ntpdate configuration file used for the ntpdate application on Ubuntu, Kubuntu, and Debian Linux desktop systems.

On Ubuntu and Kubuntu systems, the configuration information for your NTP client software is stored in the file /etc/default/ntpdate. By default, this configuration file tells the ntpdate client to look for information in the same configuration file as that used by the NTP server, which is the file /etc/ntp.conf. Unfortunately, this file is not present on client systems unless you have also installed the NTP server software on those systems, which is unnecessary. For that reason, we'll modify the /etc/default/ntpdate configuration file to ignore this file and explicitly reference your NTP server. To do so:

1. Change the line that defines the NTPDATE_USE_NTP_CONF variable to set the value of this variable to no, so that the ntpdate client will not attempt to consult the file /etc/ntp.conf for the list of time servers to use, since this file is not present on your system unless you have also installed the NTP server package on that system.
2. Add the IP address for your ntp server at the beginning of the list of time servers that is defined in the NTPSERVERS variable. This entry should be separated from the default value of ntp.ubuntu.com by a space.

After editing this file, "the right thing" will happen whenever the system (or you) bring up an Ethernet interface, but that won't synchronize your clock for you at the moment, since your Ethernet interface(s) are presumably already up. To do this, run the /usr/sbin/ntpdate-debian application manually from a command line to synchronize your system's clock with your NTP server.

TIP **If your client systems run for a long time (i.e., days or weeks), it is possible that their clocks may become skewed from the time server, because the `ntpdate` application (or, more properly, the `ntpdate-debian` script) runs only when you bring up an Ethernet interface. This means that your clock is not regularly synced with the NTP server, which can therefore pose a problem for the timestamps in your Kerberos tickets. To eliminate this potential problem, you can simply set up a regularly scheduled CRON job to sync your system's clock with the time server at a specified interval. I find daily synchronization to suffice, so I simply use the `sudo` command to execute my favorite text editor and create the file `/etc/cron.daily/ntpdate-debian` with the following contents:**

```
/usr/sbin/ntpdate-debian
```

After creating this file, make sure that it is executable by executing the following command:

```
sudo chmod 755 /etc/cron.daily/ntpdate-debian
```

Your system will now automatically synchronize with the time servers that you identified in the file `/etc/default/ntpdate`.

Setting Up a Kerberos Server

Installing the Kerberos admin server package (`krb5-admin-server`) and its dependencies on your server displays some basic configuration information and prompts you for the name or IP address of the Kerberos servers for your realm. Unfortunately, the information that it requests is insufficient to actually start and use Kerberos in your domain, so this section walks you through the steps that you'll need to perform to actually begin using Kerberos on a server.

After installing the `krb5-admin-server` and its dependencies on the system that you want to become your Kerberos server, do the following:

1. Edit the file `/etc/krb5.conf` using your favorite text editor and do the following:
 - In the `[libdefaults]` section, modify the `default_realm` entry so that it is set to the name of your realm (in uppercase letters). In my case, where I use the realm `VONHAGEN.ORG`, this entry looks like the following:

     ```
     default_realm = VONHAGEN.ORG
     ```

 - In the [realms] section, add an entry that identifies the hostnames or IP addresses of your primary Kerberos server (i.e., the system on which the primary Kerberos admin daemon is running) and your primary Key Distribution Center. The entry for this on my primary Kerberos server looks like the following:

     ```
     VONHAGEN.ORG = {
             kdc = 192.168.6.140
             admin_server = 192.168.6.140
         }
     ```

 - In the [domain_realm] section, add entries to identify the hosts that the KDC should recognize as being members of your realm. These are essentially masks for the domain(s) or subdomain(s) containing the hosts that will be contacting the Kerberos server. You should

enter two values for each domain or subdomain. For example, for hosts in the `vonhagen.org` domain, my entries in this section would look like the following:

```
vonhagen.org  = VONHAGEN.ORG
.vonhagen.org = VONHAGEN.ORG
```

- Unless you need to maintain compatibility with any existing Kerberos servers in your environment, set both the `krb4_convert` and `krb4_get_tickets` variables in the `[login]` section to false. The entry for this on my primary Kerberos server looks like the following:

```
krb4_convert = false
krb4_get_tickets = false
```

NOTE **The default `krb5.conf` file that is installed on Ubuntu and Kubuntu systems contain numerous entries for sites at CMU, MIT, Stanford, and other sites where Kerberos is popular. You should remove these before going into production with your Kerberos installation, but they initially serve as good examples for the types of customizations that you are making.**

2. Next, edit the configuration file for your Key Distribution Center, `/etc/krb5kdc/kdc.conf` to include your realm. In the realms section, change the name of the example domain to the name of your Kerberos realm, as in the following example:

```
[realms]
    VONHAGEN.ORG = {
    ...
```

3. Execute the `/usr/sbin/krb5_newrealm` command to initialize the realm, entering a master keyphrase for the realm when prompted to do so (this command runs the krb5_util command with specific options to create various configuration and data files that will be used by the Kerberos administration daemon).
4. Next, create an administrative principal so that you can do remote administration by executing the following command:

```
kadmin.local -q "addprinc admin/admin"
```

After entering your password, the administrative principal for your Kerberos domain will have been created.

5. Next, start an interactive `kadmin` session (using the admin principal and the `kadmin.local` application), and add principals for your server (the host the `kdc` and `kadmin` servers are running on) and the first client system that you will be testing it with, your user name, and the first client that you want to test with, as in the following example:

```
# kadmin.local -p admin/admin
kadmin.local: addprinc -randkey host/server.local.network
kadmin.local: addprinc -randkey host/workstation.local.network
kadmin.local: addprinc username
```

In these examples, replace *server* with the name of your server, *workstation* with the name of the client workstation, and *local.network* with the name of your realm. In my case, these would read something like kerberos.wvh.org and ubuntu.wvh.org because I tend to use kerberos as the hostname of my Kerberos servers. The add principal commands for the server and workstation principals use random keys when they are created (specified by the -randkey option) because no one will ever log in as those hosts, but you want them to exist as entities in the Kerberos database. You will be prompted for your password after creating a principal for your login—a random key won't do there, because you will actually want to be able to log in at some point.

TIP **When setting up kerberos for the first time and defining the principal for the user that you will be using for testing purposes, I suggest that you initially set the password for your user principal to something other than your standard Ubuntu or Kubuntu password. This enables you to make sure that the authentication process on your client systems is actually using Kerberos—after doing so, you can delete your existing principal (using the kadmin command's delprinc command) and then create another with your standard password.**

Before exiting the kadmin.local command, create a keytab entry for your Kerberos server. The keytab files provides default information about a host, such as the encryptions protocol that it uses (I typically use the defaults):

```
kadmin.local: ktadd host/server.local.network
```

As you did previously, replace *server* with the name of your server and *local.network* with the name of your realm. In my case, this entry would read something like kerberos.wvh.org. You can then exit from the interactive command environment provided by the kadmin.local command by issuing the quit command:

```
kadmin.local: quit
```

6. That's it for your server. As a matter of good form, I generally restart the kadmin daemon and the KDC server at this point, to verify that they will start correctly when I next boot the system. The commands to do this are the following:

```
/etc/init.d/krb5-admin-server restart
/etc/init.d/krb5-kdc restart
```

Now, off to the client to configure it to use Kerberos and your server!

Setting Up a Kerberos Client

When installing the krb5-user and libpam-krb5 packages, you will be prompted for the name of your Kerberos server and the name of your KDC server, which you can enter these as hostnames or IP addresses. Entering IP addresses is safest, because it eliminates the chance that DNS problems could prevent you from locating your Kerberos servers, but your clients obviously won't know how to find your Kerberos servers if you change their IP addresses. I typically resolve this by entering hostnames for the Kerberos and KDC servers and making sure that the IP addresses of these hostnames are listed in my /etc/hosts file. I then make sure that the /etc/nsswitch.conf file on each client uses both dns and files (in that order) in the hosts entry, which gives me both the flexibility of DNS and a fallback IP address if DNS is unavailable.

After entering these packages, use your favorite text editor to edit the file

/etc/pam.d/common-auth, adding the following line after the line that invokes pam_unix.so:

```
auth sufficient   pam_krb5.so use_first_pass forwardable
```

Make sure that this entry does not contain any typos! After saving this file, you can now test that Kerberos is working by using the kinit command, which attempts to get Kerberos tokens for the specified user (in this case, wvh):

```
kinit wvh
```

You'll be prompted for your password and should not see any errors. You can verify that this actually worked by executing the klist command, which provides information about the Kerberos tokens help by the current user, and (if all is well) will return something like the following:

```
# klist
Ticket cache: FILE:/tmp/krb5cc_1000_eYJHEN
Default principal: wvh@VONHAGEN.ORG

Valid starting     Expires            Service principal
10/22/08 05:52:52  10/22/08 15:52:52  krbtgt/VONHAGEN.ORG@VONHAGEN.ORG
renew until 10/23/08 05:52:52

Kerberos 4 ticket cache: /tmp/tkt0
klist: You have no tickets cached
```

If the kinit command returns an error or does not grant valid tickets, you can enable debugging during the Kerberos authentication process to attempt to identify (and then resolve) the problem. To activate debugging, edit the file /etc/pam.d/common-auth, using your favorite text editor and add the keyword "debug" to the end of the Kerberos entry, which should then look like the following:

```
auth   sufficient   pam_krb5.so use_first_pass forwardable debug
```

You can then execute the kinit command again and monitor /var/log/auth.log, which should contain information that will help you understand why Kerberos is not working.

Once you can successfully execute the kinit command and see tokens using the klist commands, use the ssh command to try to log in on your Kerberos server via its loopback address as a final test. If this works, you can then log out of the SSH connection, completely log out of your client, and log back in with your Kerberos password. Congratulations—you're using Kerberos!

Kerberos can be very useful as a centralized authentication mechanism, though you will have to create Kerberos principals for every user on your network (i.e., in your Kerberos realm) before they can log in on any client system that you have configured to use Kerberos authentication. As configured in this section, Kerberos requires that all of your users have local accounts (because the Kerberos authentication check in /etc/pam.d/common-auth follows the standard pam_unix.so authentication check).

Summary

This chapter started off by explaining the concepts of login users and protection groups on Ubuntu systems. It explained how Linux systems use Pluggable Authentication Modules to verify a user's identity, that a user is authorized to log in at a specific time, and perform other authentication-related tasks when logging in on an Ubuntu system. It explained the basics of file and directory protection using users and groups on Ubuntu systems, and then discussed how users perform tasks that require system privileges on Ubuntu systems using the sudo command. The last sections of the chapter explained more advanced file and system access methods—Access Control Lists, a more powerful and modern mechanism for enabling or denying access to specific files and directories, and how to install and configure Kerberos authentication on your system to provide a centralized, network-oriented authentication mechanism.

Bonus Chapter 5 discusses AppArmor and SELinux, two system-level mechanisms for guaranteeing the security of the binaries and services that you are running on your Linux system. Information about system exploits and vulnerabilities travels at the speed of the Internet, and AppArmor and SELinux provide mechanisms for helping make sure that your system is secure in the first place, and that the impact of any newly discovered and publicized vulnerabilities can be contained and minimized automatically.

Bonus Chapter 5

Managing System Policies and Security

Taking the time to implement and maintain security policies is like working out at the gym—it takes time and can be a fair amount of work, and although it generally improves the health of a system, you can never be sure which problems it's helped you avoid. System security is a generic term that spans a huge number of topics but can generally be divided into several classes: restricting access to your systems, restricting the services that your systems provide and how those services operate, and restricting what users can do on a system once they get access to it. The first two of these topics are discussed in Chapter 25, "Network Configuration and Security"; this chapter focuses on the third.

UNIX-like systems such as Linux are traditionally much more secure than their Microsoft Windows counterparts because they were designed as multiuser operating systems where different users have always had different privileges. Being conscious of the privileges required for different operations should be built into every UNIX and Linux sysadmin.

The Ubuntu and Kubuntu Linux distributions have always been security conscious. For example, Ubuntu pioneered the use of `sudo` as the primary mechanism for executing commands as a privileged user. This is a surprise to almost everyone the first time that you encounter Ubuntu and takes time to get used to, especially if you've been working with more traditional UNIX and Linux systems for a long time. To be completely honest, I probably still spend more time in a root shell than I should, using `sudo -s` as a convenient equivalent to the traditional `su` command. Old habits die hard, but at least I'm much more conscious of when and how I need to execute privileged commands nowadays.

Unfortunately, executing commands as the root user, whether in a root shell via `su` or via the `sudo` command, is the big hammer for many sysadmins. "When something doesn't work, try running it as root" is a common mantra when problems occur. When this solves a problem, some sysadmins look no further. Many UNIX and Linux commands require read access to privileged resources in order to display information about them but don't need to be able to write

IN THIS CHAPTER

Comparing AppArmor and SELinux

Configuring and Using AppArmor

Configuring and Using SELinux

to those resources. Command-line administrative commands traditionally understand this separation of functionality, but many graphical administrative commands have traditionally taken the big hammer approach.

Luckily, as graphical administrative utilities and Linux itself have matured, a number of technologies have emerged that provide more granular management of access to system resources. Ubuntu and Kubuntu systems provide some powerful, built-in security mechanisms that do a great job of separating basic and secure access requirements.

This chapter discusses two mechanisms for improved access management that are available for the Ubuntu and Kubuntu Linux distributions:

- AppArmor, which was originally developed by Immunix as Immunix Subdomain and acquired by Novell for use with its SUSE Linux distribution when Novell bought Immunix in 2005
- SELinux (Security-Enhanced Linux), which was originally developed by the National Security Agency (NSA) in the United States

Although AppArmor is installed by default (in a limited fashion) on Ubuntu and Kubuntu systems, SELinux is also available in the Ubuntu repositories used by these distributions. AppArmor and SELinux each have their own advantages, disadvantages, and legions of fanatical devotees. Which of these you use is up to you, unless you are already a party member in one of these camps or a specific system security mechanism is mandated where you work.

Overview and Comparison of AppArmor and SELinux

Both AppArmor and SELinux are examples of what are known as Mandatory Access Control (MAC) systems. At the operating system level, MAC systems limit the ability of processes or threads to access other operating system resources such as files, memory, network and other device resources, and so on. This is typically implemented by associating security attributes with each operating system object (application/process, thread, file, memory region, network or device descriptor, etc.). When a process attempts to access some other system resource, the kernel uses rules to compare the security attributes of the process and each resource that it wants to access, determining whether each access should be allowed. Because MAC systems are explicitly designed to protect system resources, MAC mechanisms typically focus on enforcing an application and resource security architecture by limiting privileges, restricting filesystem access, preventing privilege escalation, limiting resource access to localize the potential impact of security violations, and so on.

NOTE **The phrase "Big MAC Attack" is used at least once by all system administrators who are implementing a MAC-based security mechanism. It's usually best to get this out of your system early, as I've just done.**

The rules used by MAC systems are controlled by system administrators and cannot be overridden by users. This differentiates MAC systems from traditional access control mechanisms such as ACLs and UNIX user, group, and other permissions, which are known as Discretionary Access Control (DAC) systems because access to the resources "owned" by a user can be overridden at that user's discretion by changing permissions.

Another common access control mechanism is Role-Based Access Control (RBAC), which restricts user access to applications, files, and directories based on the "role" of user of a computer system. These roles correspond to the types of tasks that a user is expected to perform on a computer system, each of which requires access to different applications, files and directories, and so on. RBAC differs from access control mechanisms such as ACLs because restrictions (and privileges) are based on general, logical roles rather than user-specific permissions on system resources and filesystem objects. The most common examples of roles are normal users, who need to be able to execute many applications but generally only write to their home directories and the system's temporary directories; system administrators, who require privileged access to all resources on a system; and operators, who require privileged access to a subset of those system resources, and require a lower level of privileged access. Although AppArmor and SELinux are MAC-based security mechanisms, they can be used to simulate RBAC, as I'll explore later in this chapter.

Many other acronyms and associated security mechanisms are available for those who are seriously interested in system security. Attribute-Based Access Control (ABAC), Context-Based Access Control (CBAC), and Lattice-Based Access Control (LBAC) are interesting alternatives to MAC, DAC, and RBAC. There are probably many more, but I may not have the necessary privileges to even know about them.

Identifying Resources

The most fundamental difference between AppArmor and SELinux is that AppArmor rules, known as profiles, are path-based, while SELinux rules are based on labels associated with the file and directory inodes in a filesystem. Inode labeling is a low-level solution with obvious security benefits because security attribute labels will travel with an inode if the file associated with it is renamed or moved. On the other hand, it can only be used on filesystems that support enhanced attributes, which immediately rules out standard networked filesystems such as NFS.

Labeling each file and directory inode in today's huge filesystems can take a substantial amount of time, but you really only have to label executables and associated libraries in the directories used by the system and by application. However, you should configure your system to scan for and report new, unlabeled files of these types in these directories. An operating system update or upgrade that installs new executables and shared libraries will require relabeling if you are using SELinux, while the same upgrade would be largely transparent if you were using AppArmor.

An additional, label-related issue is the fact that many backup solutions don't save extended attributes (including standard utilities such as `tar`). If this is the case with the backup utility that you are using, you may lose your labels if you are forced to restore all or a part of a given filesystem from backup media.

The profiles used by AppArmor to monitor and control executables and the files associated with them are text files that are stored in the directory `/etc/apparmor.d`. To simplify profile creation and maintenance, these profiles load general system settings and service definitions from subdirectories of `/etc/apparmor.d`:

- `tunables`: Contains files that provide system-wide information such as the location of user home directories and the mount point for the `procfs` filesystem that is used to store and monitor information about running processes on the system
- `abstractions`: Contains files that provide information about specific files and directories that are associated with logical services that span multiple processes, such as authentication, `X` (for X Window system applications), `perl`, and `python` for scripts that require these languages, and so on

SELinux configuration files are located in the directory `/etc/selinux` but are not installed on Ubuntu or Kubuntu systems by default because AppArmor is the default system security mechanism. SELinux policies are installed in subdirectories of the `/etc/selinux` directory with the name of the policy that you are using as defined in the configuration file `/etc/selinux/config`.

Rule Readability and Ease of Creation

AppArmor rules are written in a somewhat readable text format, while SELinux rules are written in a macro language and must be subsequently processed by the M4 macro processor for actual use. AppArmor's rule specification language uses familiar syntax and notation such as shell syntax wildcards, rwx to denote read, write, and execute permissions, and so on. Most people see the readability of AppArmor rules as a distinct advantage for readability and auditability reasons, but it unfortunately makes it more difficult to develop graphical tools.

Comparison Summary

A good general way of comparing and contrasting AppArmor and SELinux is to think of the AppArmor as focusing on data integrity, while SELinux focuses on data security. While there is substantial overlap between these two goals, AppArmor focuses on protecting the system itself from invasion and damage, while SELinux focuses on protecting the data on the system from being accessed by unauthorized users or applications.

As things stand at the moment, the differences between the two MAC systems generally boil down to:

- The tradeoff between configuration, management, and maintenance time and the level of security that you need
- The existence of an established base of rules, domain definitions, and policies that you can easily adapt to new systems
- The tools that are provided for MAC configuration, management, and maintenance

In general, I'm an AppArmor fan. I have never needed the completely anal security that SELinux provides, and I doubt that the CIA, NSA, or any other similar TLA organization is looking to hire me at any point in the future. Although there aren't any graphical tools for AppArmor policy development that I'm aware of, I'm OK with that. It's not like my grandfather is going to be creating lots of AppArmor profiles.

In general, I find AppArmor's profiles and their contents easy to work with—they are readable and can easily be modified using my favorite software package, `emacs`. Because AppArmor is path-oriented, new profiles can easily be tested and deployed without requiring filesystem labeling and relabeling. AppArmor's learning mode makes it relatively easy to create profiles for new applications, especially when no precedent exists, by recognizing violations and then using the log of those violations as the basis of a new profile. To date, learning mode has always worked well for me in these cases (as illustrated in the section entitled "Creating a Sample AppArmor Profile"), and cloning an existing profile and fine-tuning it to meet my needs is equally easy when a similar profile is already in use.

Configuring and Using AppArmor

AppArmor is the default MAC system that is used on Ubuntu and Kubuntu systems, and controls access to system resources through profiles that define the files, directories, and level of access that different applications have to different resources. These path-based profiles make it easy to examine the

restrictions that are placed on various system executables and associated files, modify them if necessary, and create new profiles for site-specific applications.

Although I'll go into more detail about how AppArmor profiles work later in this section, it's important to understand a few basic concepts before proceeding. AppArmor profiles control access to system resources in two basic modes:

- **complain mode:** Profile violations are detected and logged, but are allowed to occur.
- **enforce mode:** Profile violations are detected, logged, and prevented.

The default mode of an AppArmor profile is enforce mode but, as you'll see in the next section, most AppArmor profiles are initially loaded in complain mode so that your (or other system administrators) have the opportunity to see them in action (via log messages) and fine-tune them without effectively denying access to any process that is protected by a profile.

The next few sections explain how and where AppArmor, its configuration information, and the AppArmor profiles are installed on your system and explore the format of these profiles to help you create and manage access to the resources on your system.

Installing and Configuring AppArmor

AppArmor is installed by default on Ubuntu and Kubuntu systems and is automatically integrated into the system boot and shutdown processes through symlinks in the `/etc/rcS.d` directory. The startup scripts in this directory are executed at system startup time, before the scripts for any specific runlevel are executed. The symbolic link for starting AppArmor is named `/etc/rcS.d/S37apparmor` and runs immediately after all of the local filesystems that will be used by the system have been mounted (and also before any network filesystems have been mounted). This guarantees that AppArmor profiles can be used to protect any binary in any mounted filesystem with the exception of the `init` process itself. As you might expect, AppArmor profiles can only protect binaries if the profiles that apply to a given binary have been loaded before the binary is executed.

TIP **For general information about the location and organization of AppArmor profiles, see the "Identifying Resources" section, earlier in this chapter. For detailed information about the content and structure of an AppArmor profile, see the next section, "AppArmor Profiles and Profile Syntax."**

As installed by default on Ubuntu and Kubuntu systems, AppArmor only installs and activates three profiles, although it installs many general service definitions in the `/etc/apparmor.d/abstractions` directory that you can use when creating or installing additional profiles. The following command and related output shows the status of the default AppArmor installation on a Ubuntu or Kubuntu system:

```
$ sudo aa-status
apparmor module is loaded.
3 profiles are loaded.
3 profiles are in enforce mode.
   /usr/share/gdm/guest-session/Xsession
   /usr/sbin/cupsd
   /usr/lib/cups/backend/cups-pdf
0 profiles are in complain mode.
1 processes have profiles defined.
1 processes are in enforce mode :
   /usr/sbin/cupsd (4876)
```

```
0 processes are in complain mode.
0 processes are unconfined but have a profile defined.
```

AppArmor's default configuration and associated profiles are designed to protect the standard system processes that most systems would be running after a default Ubuntu or Kubuntu installation, and which might be accessible to other systems. These are the CUPS print daemon (`cupsd`, which remote systems could access), guest login sessions via the GNOME display manager (`Xsession`), and PDF generation via the CUPS print server (`cups-pdf`).

AppArmor's default configuration is intended to be lightweight but helpful, but is analogous to having the world's best home security system and only using it to protect your garage. Don't worry—AppArmor also provides default profiles for many other processes, but these profiles are simply not installed by default on your Ubuntu or Kubuntu system. For those of us who install or simply use other services, such as Samba, a time server, the Avahi network resource discovery daemon, advanced DNS-related daemons, and so on. The Ubuntu repositories provide the `apparmor-profiles` package, which you should probably install on your system.

After installing the `apparmor-profiles` package on your system using your favorite package management software (as discussed in Chapter 23, "Adding, Removing, and Updating Software"), you'll need to restart your system so that your new AppArmor profiles can be loaded before the processes that they are associated with, as explained earlier in this section. After restarting your system (or manually restarting every application or daemon for which a profile exists, which is much trickier), you'll see that a lot more is going on (and protected) on your system. The following command and related output shows the status of the updated AppArmor installation on a Ubuntu or Kubuntu system:

```
$ sudo aa-status
apparmor module is loaded.
16 profiles are loaded.
4 profiles are in enforce mode.
   /usr/share/gdm/guest-session/Xsession
   /usr/lib/cups/backend/cups-pdf
   /usr/sbin/cupsd
   /usr/sbin/avahi-daemon
12 profiles are in complain mode.
   /usr/sbin/identd
   /usr/sbin/ntpd
   /sbin/klogd
   /usr/sbin/dnsmasq
   /usr/sbin/nmbd
   /sbin/syslogd
   /usr/sbin/smbd
   /sbin/syslog-ng
   /usr/sbin/traceroute
   /usr/sbin/nscd
   /bin/ping
   /usr/sbin/mdnsd
9 processes have profiles defined.
3 processes are in enforce mode :
   /usr/sbin/avahi-daemon (6787)
   /usr/sbin/cupsd (6880)
   /usr/sbin/avahi-daemon (6786)
6 processes are in complain mode.
```

```
   /usr/sbin/nmbd (7111)
   /usr/sbin/smbd (7177)
   /sbin/klogd (6741)
   /sbin/syslogd (6688)
   /usr/sbin/smbd (7113)
   /usr/sbin/ntpd (7942)
0 processes are unconfined but have a profile defined.
```

Now that's more like it! However, while this looks more secure, I haven't really explained what profiles look like and how they work. The next section fixes that problem.

NOTE **If you install the `apparmor-profiles` package but have not rebooted your system or restarted every profiled daemon/application after doing so, the last section of your output will look like the following:**

```
8 processes are unconfined but have a profile defined.
  /usr/sbin/smbd (5154)
  /sbin/syslogd (4733)
  /usr/sbin/ntpd (6002)
  /usr/sbin/avahi-daemon (4830)
  /usr/sbin/avahi-daemon (4831)
  /usr/sbin/nmbd (5152)
  /usr/sbin/smbd (5206)
  /sbin/klogd (4785)
```

This is because the applications were already running when their profiles were loaded. Trust me—you need to restart those processes or reboot your system. (Or, as we say in Pennsylvania, "Your system needs rebooted.")

AppArmor Profiles and Profile Syntax

AppArmor's profiles are located in the directory `/etc/apparmor.d` on Ubuntu or Kubuntu systems, and have names of the form *`full.path.to.application`*. For example, the name of the profile for the `/bin/ping` application is `/etc/apparmor.d/bin.ping`.

AppArmor's profiles are simple text files, where "simple" means that they are ASCII text files that contain a limited number of different types of statements, not that you're guaranteed to be able to look at one and think, "Oh yeah, I see what that does," in 10 seconds or so. Like the configuration data for most well-designed applications, the basic syntax of an AppArmor configuration file makes it possible to identify and share security information that is common across multiple applications, while still enabling you to define application-specific restrictions. The general structure of an AppArmor profile is the following:

```
# comment
application-specific-variable definition
 #include <file/containing/variable/definitions>

/full/path/to/application flags=(FLAG) {
  #include <file/containing/other/rules>

  network DOMAIN ACCESS-TYPE,

  capability POSIX-CAPABILITY,
```

```
    /full/path/to/file/or/directory PERMISSIONS,
    /full/path/to/other/file/or/directory PERMISSIONS,

    ^OPTIONAL-SUBPROFILE {
      /full/path/to/related file PERMISSIONS,
      /full/path/to/other/related/file PERMISSIONS,
    }
}
```

Each of these types of entries has its own syntax, which are the following:

- `# comment`: Lines that contain the hash mark (#) character in the first column are comments that have no impact on the profile but can have significant impact on anyone's understanding of it. Most of the standard AppArmor profiles include comments that are copyright notices, last-modified timestamps, or hints for editing the file in an editor such as vi (vim). This is not really very helpful unless you're a lawyer (or a vi user, which is at least more credible than being a lawyer). If you create your own profiles, I encourage you to do better at commenting your profiles unless you actually hate all of your coworkers or are a devotee of the "security though obscurity" model that has plagued UNIX and Linux systems for decades.
- `application-specific-variable definition`: AppArmor variable definitions enclose the name of the variable that you are defining in within curly brackets preceded by an @ symbol followed by an equals sign (=), and a space-separated list that defines the value of that variable. A common example is a variable that defines the home directories for the users on your system, as in the following:

  ```
  @{HOME} = /home/*/ /root/
  ```

- Variable definitions can use other variable definitions. For example, since you may want to handle the root user's home directory differently than those of "normal" users in a rule, you might use the following definitions:

  ```
  @{USERHOMES} = /home/*/
  @{ROOTHOME} = /root/
  @{HOMES} = @{USERHOMES} @{ROOTHOME}
  ```

- `#include <`*`file/containing/something`*`>`: The path to a file containing variable definitions or rules. The pathname of the file is enclosed within angle brackets and is relative to AppArmor's profile directory, which is the directory `/etc/apparmor.d` on Ubuntu or Kubuntu systems. The `#include` statement can be used to include the contents of file containing variable definitions, capability information, or rules. The type of information that the included file contains must be syntactically valid at the point in the profile where the `#include` statement appears. For example, you cannot use `#include` to include AppArmor rules outside of the application section of a profile file.
- */full/path/to/application flags=(complain) { ... }*: The full path to the application that this profile applies to, followed by an option flags setting and curly brackets that include the rules, network, and capability definitions for that application. The optional `flags=(complain)` entry enables you to specify the mode in which you want this profile to operate, which enables you to create profiles that operate in a mode other than enforce mode, which is the default mode for AppArmor profiles. If this statement is present in an AppArmor profile, it therefore usually says `flags=(complain)`.

- Note that each line in the rules, network, and capabilities section ends with a comma, including the last such entry.
- *network DOMAIN ACCESS_TYPE,*: Enables you to identify a specific network domain and the type of access that an application requires. Supported *DOMAIN*s include `appletalk`, `bluetooth`, `inet`, `inet6`, `ipx`, and `netrom`, though most of us will primarily care about `inet` (an IPv4 Ethernet network) and `inet6` (an IPv6 Ethernet network) networks. Supported *ACCESS_TYPE* values include `dgram`, `stream`, and `seqpacket`, although most of us will primarily care about `dgram` (Datagrams) and `stream` access.
- `capability` *POSIX-CAPABILITY,*: Enables you to identify specific POSIX capabilities that the application requires. Examples include `net_bind_service` (for applications that need to bind to privileged ports), `setgid` and `setuid` (for applications that change their group or user ID when they are started), `sys_chroot` (for applications that change their root directories to limit the portions of a system that they can access), and so on. The complete list of POSIX capabilities can be inferred from the capabilities man page (in section 7 of the Linux manual), which Ubuntu and Kubuntu systems do not provide. See `www.kernel.org/doc/man-pages/online/pages/man7/capabilities.7.html`.
- */full/path/to/file/or/directory PERMISSIONS,*: Provides access or execute permissions for the specified file or directory. A *PERMISSIONS* specification can be composed of two types of values:
 - **Access:** Defines the type(s) of access that an application requires for a file or directory. Possible values are `a` (append), `k` (locking), `l` (linking), `r` (read), and `w` (write). Appending to a file and writing to a file are mutually exclusive because they are different write models. Writing to a file destroys any previous contents of that file, while appending simply adds to the end of the file. For example, you always append to log files, while you typically simply write to data files.
 - **Execute:** Defines the execute mode for a specific executable application that can be called by the application to which the current profile applies. Possible values are `ix` (same as the "Parent" application), `m` (enables data to be mapped into memory and executed), `Px` (requires that a separate profile exist for the specified application, which does not inherit any variables, network settings, or capability information from the "parent" application's profile), `px` (requires that a separate profile exist for the specified application, which inherits the variables, network settings, and capability information from the "parent" application's profile), `Ux` (executes the specified application without a profile, and does not inherit any variables, network settings, or capability information from the "parent" application's profile), and `ux` (executes the specified application without a profile, though the specified file inherits the variables, network settings, and capability information from the "parent" application's profile). The `ux` and `Ux` execution modes essentially remove any execution limitations for the specified application and are often used when developing a new profile. You should not use these modes in deployed profiles for obvious reasons.

NOTE **With AppArmor 2.3, file/directory rules can be general or owner-specific. Owner-specific rules are identical to regular file/directory rules, but begin with the owner keyword. Such rules are conditional upon the user being the owner of the file. If an owner's rules overlap with a regular file/directory rule, the latter is used because it specifies permissions regardless of who owns the file.**

- *^OPTIONAL-SUBPROFILE { ... }*: The up-arrow character identifies the name of a subprofile, which is known in AppArmor-speak as a "hat," and is followed by curly brackets that enclose rules that are specific to subprocesses of the application to which the general profile applies. Subprofiles are typically only required in applications that execute and run many other applications, some of which have distinct resource requirements that may require different permissions than the main application. Subprofiles are most commonly used for Web-related application, specifically the Apache (`httpd`) and Apache2 (`apache2`) Web servers. This is not a book on AppArmor, so I don't have space to discuss all of the nuances of subprofiles, especially for complex applications such as Apache2. However, it isn't surprising that the Novell folks have some great docs on subprofiles and Apache. See the section on "Additional Information about AppArmor" for a link to the Novell *AppArmor Administration Guide*, which is a great resource. Novell has always had great documentation, even before they cared about Linux, and they're still true to that ideal. You'll have to ignore all the YaST stuff (the SUSE/Novell administration tool), but it's much better to be able to ignore something than it is to not have any information in the first place.

The previous list contained a lot of general information, so how about an actual example? The following is the AppArmor profile for the `/bin/ping` application:

```
#include <tunables/global>
/bin/ping flags=(complain) {
  #include <abstractions/base>
  #include <abstractions/consoles>
  #include <abstractions/nameservice>

  capability net_raw,
  capability setuid,
  network inet raw,

  /bin/ping mixr,
  /etc/modules.conf r,
}
```

I'm not going to expand each included file (you can do that in the privacy of your own Ubuntu or Kubuntu system), but you can see that this profile is loaded in complain mode and that the capabilities, network statement, and rules in this file have the following meanings:

- `capability net_raw`: Requires raw network access.
- capabilities `setuid`: Requires the ability to set the user ID upon execution to the owner of the file, which is the root user. This enables ping users to send raw packets over the network.
- `network inet raw`: This requires raw access to IPv4 networks, which enables the application to craft its own packets and send them over an IPv4 network.
- `/bin/ping mixr`: The specific binary for this application can be read, mapped into memory, and executed. The `ix` statement seems somewhat redundant but ensures that the specific binary has the default capabilities and types of network access that are defined in the profile.
- `/etc/modules.conf r`: The application can read the `/etc/modules.conf` file to identify network device aliases and related information.

To be honest, the profile for the `/bin/ping` application is the simplest one that is installed on an Ubuntu or Kubuntu system by the `apparmor-profiles` package. Although I actually don't like saying "Use the source, Luke" as a technique for learning how application code works, I don't feel so bad about saying "Read the configuration files, Luke" as a mechanism for learning how to create or modify your own AppArmor profiles.

As an assignment, I suggest that you now use your favorite text editor to create an AppArmor profile for the Java virtual machine and all of the `jar` files on your system. Just kidding! The next section explains the utilities that AppArmor provides to create and work with your profiles and manage your AppArmor installation.

Using AppArmor Utilities

Installing AppArmor also installs a number of related utilities on your system. These utilities, which make it easy (or relatively easy) for you to create profiles, change profile states, show AppArmor status, and perform other AppArmor-related tasks, are the following:

- `aa-audit`: This switches an existing AppArmor profile into audit mode, in which security policy is enforced and all access to any file that is listed in the profile is logged to the system log regardless of whether that access was successful or not.
- `aa-autodep`: This generates a minimal AppArmor profile for executables (both scripts and binaries) and any libraries that they use. The resulting profile contains an entry for at least the specified file and an `include` directive that includes the basic profile entries that are required for most programs.
- `aa-complain`: This switches an existing AppArmor profile into complain mode, in which security policy is not enforced but all access to any file that is listed in the profile is logged to the system log regardless of whether that access was successful or not.
- `aa-enforce`: This switches an existing AppArmor profile into enforce mode, in which security policy is enforced and all access to any file that is listed in the profile is logged to the system log regardless of whether that access was successful or not. This is the default mode for an AppArmor profile.
- `aa-genprof`: This creates a profile for a specified application if one does not exist. It then switches the new or existing profile into complain mode, writes a mark to the system log, and tells you to start the specified application in another window so that it can monitor its behavior. You are then presented with two choices: enter S to scan the system log for messages related to that application and add the associated file to the profile, or enter F to set the new profile to enforce mode and exit. If you request scanning (S), `aa-genprof` will use `aa-logprof` to identify any violations of the existing policy, update the policy, and repeat this process until you break this cycle by entering F to finish the process. Once you enter F, `aa-genprof` will set the profile(s) associated with the application into enforce mode and will exit.
- `aa-logprof`: Used when creating profiles, this utility extracts and reviews information about AppArmor profiles that are running in complain mode, and uses that information when creating a new profile.
- `aa-status`: This displays status information about AppArmor on your system, showing the application profiles that are active on your system, the modes in which they are operating, and any discrepancies between available profiles and running processes. Sample output from the `aa-status` command was shown at the beginning of the section on "Installing and Configuring AppArmor."

- `aa-unconfined`: This lists all processes on your system that do not have an AppArmor profile loaded but that are using a TCP or UDP port on your system. By providing a list of applications for which you may need to create AppArmor profiles, this utility can help you pass the time when all of your other system administration duties are completed.

It's easy to say that the applications provided with AppArmor will help create profiles for new applications, but as they say on TV, "the proof is in the profile." The next section shows how to use these applications to create a sample profile for a new system application.

Creating a Sample AppArmor Profile

I'm a big fan of Dropbear (`http://matt.ucc.asn.au/dropbear/dropbear.html`), which is a lightweight alternative to the standard OpenSSH server. I'm not sure whether this is because I've done a lot of work in the embedded Linux space or simply because I don't like to needlessly squander system memory, but Dropbear is my SSH server of choice. Unfortunately, AppArmor doesn't include a profile for Dropbear (or the OpenSSH server, for that matter). Dropbear is a system service that runs a daemon to respond to incoming network requests and offers a secure, encrypted shell connection when used in its most standard fashion. All of which combine to make it an excellent test case to see how easy it is to use the standard AppArmor utilities.

The first step in creating a profile using the AppArmor utilities is to shut down the application for which you are creating a profile. This enables the AppArmor utilities to monitor the startup and standard runtime behavior of the application. To shut down Dropbear, I used the following command:

```
sudo /etc/init.d/dropbear stop
```

Next, execute the `aa-genprof` command with the full pathname of the associated application as an argument. This command and its initial output are the following:

```
$ sudo aa-genprof /usr/sbin/dropbear
Connecting to repository.....

Writing updated profile for /usr/sbin/dropbear.
Setting /usr/sbin/dropbear to complain mode.

Please start the application to be profiled in
another window and exercise its functionality now.

Once completed, select the "Scan" button below in
order to scan the system logs for AppArmor events.

For each AppArmor event, you will be given the
opportunity to choose whether the access should be
allowed or denied.

Profiling: /usr/sbin/dropbear
[(S)can system log for SubDomain events] / (F)inish
```

At this point, I opened another tab in my terminal session and executed the following Dropbear startup command:

```
sudo /etc/init.d/dropbear start
```

To check that AppArmor complain-mode messages are being written to the system log, use a command like the following:

```
$ sudo grep dropbear /var/log/messages
Dec 14 14:26:06 ubuntu kernel: [12245.881212] type=1502 \
  audit(1229801166.910:574): operation="socket_sendmsg" \
  family="inet6" sock_type="stream" protocol=6 pid=11134 \
  profile="/usr/sbin/dropbear"
Dec 14 14:26:06 ubuntu kernel: [12245.881266] type=1502 \
  audit(1229801166.910:575): operation="file_permission" \
  requested_mask="w::" denied_mask="w::" fsuid=1000 \
  name="/dev/pts/2" pid=11135 profile="null-complain-profile"
Dec 14 14:26:06 ubuntu kernel: [12245.881309] type=1502
  audit(1229801166.910:576): operation="socket_sendmsg" \
  family="inet6" sock_type="stream" protocol=6 pid=11134 \
  profile="/usr/sbin/dropbear"
```

You can then return to the window or terminal session tab in which you started the `aa-genprof` utility and enter S to tell the `aa-genprof` utility to scan the log for related messages. The `aa-genprof` command displays the following output:

```
Reading log entries from /var/log/messages.
Updating AppArmor profiles in /etc/apparmor.d.
Profiling: /usr/sbin/dropbear
[(S)can system log for SubDomain events] / (F)inish
```

Repeat this step a few times, and then respond F to tell the `aa-genprof` utility to complete the profile creation process. You will see output like the following:

```
Setting /usr/sbin/dropbear to enforce mode.
Reloaded SubDomain profiles in enforce mode.
Finished generating profile for /usr/sbin/dropbear.
```

Since we're just creating a new profile, it may be incomplete and, therefore, should not be running in enforce mode, which will prevent access to the application if the application tries to access resources that are not listed in the profile file. You can use the `aa-complain` command to change it to complain mode, as in the following example, which shows the command and its output:

```
$ sudo aa-complain /usr/sbin/dropbear
Setting /usr/sbin/dropbear to complain mode
```

You can then use the `cat` command to display the profile that has been created, as in the following example:

```
$ cat usr.sbin.dropbear # Last Modified: Sat Dec 11 14:25:15 2008
#include <tunables/global>

/usr/sbin/dropbear flags=(complain) {
  #include <abstractions/base>
}
```

This shows that the rule is indeed in complain mode and that some rules have been created, although none of the file-specific log messages that you saw earlier seem to be present (more about that later).

We can assume that `aa-genprof` ran `autodep` to identify the dependencies in the `/usr/sbin/dropbear` binary, all of which are presumably covered in the file `/etc/apparmor.d/abstractions/base`. To verify that this is correct, use the `ldd` command to manually check for the dependencies in the `/usr/sbin/dropbear` executable, as in the following example:

```
$ ldd /usr/sbin/dropbear
linux-vdso.so.1 =>  (0x00007fff74bfe000)
libutil.so.1 => /lib/libutil.so.1 (0x00007fe66c5f0000)
libz.so.1 => /usr/lib/libz.so.1 (0x00007fe66c3d8000)
libcrypt.so.1 => /lib/libcrypt.so.1 (0x00007fe66c1a0000)
libc.so.6 => /lib/libc.so.6 (0x00007fe66be2e000)
/lib64/ld-linux-x86-64.so.2 (0x00007fe66c7f3000)
```

The first library listed is a virtual dynamic shared object of the Linux kernel, but the others should be matched by rules in the `#included` rules file. You can then examine the file `/etc/apparmour.d/abstractions/base` to see if it contains entries that cover the specified libraries, which it does. These rules are the following:

```
  /lib/lib*.so*                  mr,
  /usr/lib/*.so*                 mr,
  /lib64/ld-*.so                 mrix,
```

As mentioned earlier, it is odd that the sample profile doesn't contain entries for the log messages that we saw earlier. However, these can be added by manually running the `aa-logprof` program. No options are necessary when running this command, as it simply scans the system log since the last mark event, looking for new audit messages, as in the following example:

```
$ sudo aa-logprof
Reading log entries from /var/log/messages.
Updating AppArmor profiles in /etc/apparmor.d.
Complain-mode changes:

Profile:  /usr/sbin/dropbear
Path:     /etc/dropbear/dropbear_dss_host_key
Mode:     owner r
 Severity: unknown
  [1 - /etc/dropbear/dropbear_dss_host_key]

[(A)llow] / (D)eny / (G)lob / Glob w/(E)xt / (N)ew / \
 Abo(r)t / (F)inish / (O)pts
```

Your choices for handling the file identified in the `Path` statement are displayed as a numbered list of options. The default choice from this list is displayed inside square brackets—in this case, there is only one option—add a rule allowing read access to the specified file.

At this point, you must enter a character specifying how to handle the message that has been detected and the file(s) with which it is associated. The command associated with the default option in the list is enclosed within square brackets. The most common choices are to allow (A, the default in this case, since it is enclosed within square brackets) or deny (D) access to the file. Other choices are to abort (r) without saving changes to the profile, or finish (F), writing changes to the profile.

In a case such as the preceding example, where a single file is identified, it is clearly related to the application, and there is only one option for handling that file to choose from, enter A to accept this entry. The

specified rule is added to the profile, and the `aa-logprof` application displays a message like the following, which explains that action that it performed:

```
Adding /etc/dropbear/dropbear_dss_host_key owner r to profile.
```

The next log message present in `/var/log/syslog` is processed, and the relevant information about that message is displayed, as shown in the following example:

```
Profile:  /usr/sbin/dropbear
Path:     /etc/gai.conf
Mode:     owner r
Severity: unknown

  1 - #include <abstractions/nameservice>
 [2 - /etc/gai.conf]

[(A)llow] / (D)eny / (G)lob / Glob w/(E)xt / (N)ew / \
  Abo(r)t / (F)inish / (O)pts
```

In this case, there are two options:

- Use the `#include` statement to include a file that presumably contains rules for using the nameservice, which would include a rule for the file `/etc/gai.conf` (which is used to create a structure for holding network client or server socket connection information).
- Explicitly allow read access to the file `/etc/gai.conf`.

The second option is the default and seems like a reasonable rule, so you can simply press A to accept this rule. The specified rule is added to the profile, and the `aa-logprof` application displays a message like the following, which explains that action that it performed:

```
Adding /etc/gai.conf owner r to profile.
```

The next log message present in `/var/log/syslog` is processed, and the relevant information about that message is displayed. This process repeats until all new log messages have been processed, which may affect multiple profiles. When the last log message is displayed, a final message like the following is displayed:

```
= Changed Local Profiles =

The following local profiles were changed.  \
   Would you like to save them?

 [1 - /usr/sbin/cupsd]
  2 - /usr/sbin/dropbear
  3 - /usr/sbin/smbd

(S)ave Changes / [(V)iew Changes] / Abo(r)t
Writing updated profile for /usr/sbin/cupsd.
Writing updated profile for /usr/sbin/dropbear.
Writing updated profile for /usr/sbin/smbd.
```

At this point, you can enter S to save the updated entries. The specified profile files are updated and reloaded. You can then use the `cat` command to display the updated profile, as in the following example:

```
$ cat /etc/apparmor.d/usr.sbin.dropbear
# Last Modified: Sat Dec 14 15:08:34 2008
#include <tunables/global>

/usr/sbin/dropbear {
  #include <abstractions/base>

  network inet dgram,
  network inet6 dgram,
  network inet6 stream,

  owner /etc/dropbear/dropbear_dss_host_key r,
  owner /etc/dropbear/dropbear_rsa_host_key r,
  owner /etc/gai.conf r,
  owner /etc/host.conf r,
  owner /etc/hosts r,
  owner /etc/nsswitch.conf r,
  owner /etc/resolvconf/run/resolv.conf r,
  owner /etc/shadow r,
  owner /etc/shells r,
  owner /usr/sbin/dropbear r,
  owner /var/run/dropbear.pid w,
}
```

That's more like it! Since the profile is now in complain mode, you can run it for a while, collect any additional log messages that may appear, and eventually rerun the `aa-logprof` command to add rules to cover the new occurrences. You may also want to add a meaningful comment to the beginning of the file. In this case, I added the following comment:

```
# This is an AppArmor profile for dropbear, a
# light-weight alternative to the openssh-server
# package. This profile was created by running
# aa-genprof amd then subsequently using aa-logprof
# to add network and file access statements by
# scanning the log. Created by wvh@vonhagen.org
# on 11 Dec 2008.
```

This is both thoughtful and useful to the next administrator who is curious about the nonstandard profile file found on your system.

Emulating RBAC with AppArmor

The profile that I created in the previous section is a standard AppArmor profile that uses AppArmor as a mandatory access control system. However, AppArmor profiles can easily be created that support the equivalents of other access control mechanisms. The easiest of these to support is role-based access controls (RBAC).

For information about creating RBAC-like profiles, see the AppArmor FAQ (`http://developer.novell.com/wiki/index.php/Apparmor_FAQ`) for a great example that is provided in the "Can AppArmor be used to create restricted login shells?" section.

Some other sources of good information about RBAC and AppArmor are the following:

- `http://developer.novell.com/wiki/index.php/Apparmor_RBAC_in_version_2.3`
- `http://labs.novell.de/wiki/index.php/Apparmor_RBAC_in_version_2.4`

Additional Information about AppArmor

The previous sections provided a great deal of information about AppArmor but couldn't explain everything—after all, this isn't a book about AppArmor. For more information about AppArmor, see the following resources. Not all of this information is specific to using AppArmor on Ubuntu or Kubuntu systems, but it is all useful:

- `http://fridge.ubuntu.com/node/1175`: An article about AppArmor from the Ubuntu Fridge
- `https://help.ubuntu.com/community/AppArmor`: The Ubuntu Community documentation section on AppArmor
- `https://wiki.ubuntu.com/AppArmor`: Information about AppArmor on the Ubuntu wiki
- `http://en.wikipedia.org/wiki/AppArmor`: The Wikipedia article on AppArmor
- `http://en.opensuse.org/Apparmor`: The OpenSUSE page for AppArmor. Information about specific versions of AppArmor (2.3 is the latest at the time that this book was updated) is available at `http://en.opensuse.org/AppArmor/Changes_AppArmor_2_3`.
- `www.novell.com/documentation/apparmor/index.html`: The Novell AppArmor Administration Guide, which is a fantastic resource regardless of which distribution you're running
- `www.mpipks-dresden.mpg.de/~mueller/docs/suse10.3/opensuse-manual_de-pdf/APPARMOR-admin_de.pdf`: The OpenSUSE AppArmor Documentation for version OpenSUSE 10.3. This is slightly out of date, but is still an excellent source of information regardless of the Linux distribution on which you're running AppArmor.

As mentioned previously (and as I think I've shown), I find AppArmor to be a very powerful yet eminently understandable and usable system security tool. I hope that you agree.

Configuring and Using SELinux

Finally, something from the security sections of the U.S. government that does something other than read my email, spy on my network communications, and tap my phone. Who would have thought that they would come up with something that was useful, and would then actually let people use it? Go spooks! It's also comforting to know that SELinux is completely open source and that zillions of hackers have, therefore, read the code to ensure that using it doesn't add a network backdoor on each and every Linux system. Hey, is that a black helicopter circling my house?

More seriously, SELinux is an excellent and extremely powerful set of system security and associated utilities for Linux. This section is not as exhaustive as the previous section on AppArmor, for several good reasons:

- There are several good books available on SELinux, and I couldn't hope to expand much on those in the context of a chapter of this book.
- Support for SELinux on Ubuntu and Kubuntu systems is nowhere near as mature as support for AppArmor. There are many excellent developers working on both projects, and SELinux support

improves with each Ubuntu release. However, at the time that this book was updated, the best summary description for SELinux on Ubuntu and Kubuntu 8.10 systems is "broken."

TIP **If you must really run SELinux, you might want to consider running another distribution, such as CentOS or Fedora, that focuses on SELinux rather than AppArmor. I realize that this statement may seem to be heretical in the context of the "Ubuntu Bible," but my primary responsibility is to give you good advice, not simply to serve as a nerdy apostle for Ubuntu/Kubuntu. I prefer it when these two jobs are the same thing, but SELinux is that rare case in which they are disjoint.**

The following sections provide an overview of installing and configuring SELinux on your Ubuntu or Kubuntu 8.10 system, discuss the most commonly used tools that are available for working with SELinux on Ubuntu and Kubuntu systems, and identify sources of more detailed information on using SELinux on Ubuntu, Kubuntu, and other Linux distributions.

Installing SELinux

Installing SELinux on your Ubuntu or Kubuntu is relative easy in theory. To do so, simply execute the following command:

```
sudo apt-get install selinux-policy-default
```

This will remove AppArmor, install SELinux, and install a default policy for your system. Now the fun begins:

1. Change directory to `/etc/initramfs-tools/scripts/init-bottom`, and execute the following command to see how SELinux tries to load your default policy when running from the initial RAM disk:

   ```
   grep bin/load_policy _load_policy
   ```

 If the path to the `load_policy` binary is not `/usr/sbin/load_policy`, edit the file and correct the path. You must then rebuild the initial RAM disk for your kernel. You can determine which initial RAM disk your system is using by executing the following command:

   ```
   uname -r
   ```

 To rebuild your system's initial RAM disk to be able to load your SELinux policy, execute the following command, where *`STRING`* is the value returned by the previous command:

   ```
   sudo update-initramfs -u initrd.img-STRING
   ```

2. Check the file `/etc/selinux/config`. The `SELINUX` variable should be set to "permissive," and the `SELINUXTYPE` variable should be set to "default." Use the `sudo` command and your favorite text editor to correct these values if necessary.
3. Check the file `/boot/grub/menu.lst` to make sure that the kernel for your system is set to disable AppArmor and enable SELinux. You can use the value returned by the `uname -r` command in the previous step to find the entries for your current kernel, as in the following command:

   ```
   grep STRING /boot/grub/menu.lst
   ```

This command will display at least two sets of title/kernel/`initrd` entries. The line for the kernel that does not follow the title entry containing the string "(recovery mode)" should contain the following entry:

```
apparmor.enabled=0 selinux=1
```

If the kernel entry does not contain these variable settings, use the `sudo` command and your favorite text editor to add them to the end of that kernel entry. Do not add them to the kernel entry following the title that contains the string "(recovery mode)"—you'll need that to boot your system if and when you tire of experimenting with SELinux.



When trying SELinux for the first time, you should make sure that the file `/boot/grub/menu.lst` contains a reasonable `timeout` value (I suggest 5 seconds), and you may also want to comment out the `hiddenmenu` option. This will cause your system to display the GRUB menu for 5 seconds when it boots, which will give you time to select the "(recovery mode)" entry for your kernel if SELinux causes problems on your system.

At this point, you can reboot your system to try out SELinux. If things have been fixed since this book has been updated, you'll be able to log in on your system from the console. If not, your console will not accept input from your keyboard and your mouse will not be responsive. Now that's secure! If this is the case, you can press Alt+F2 to display a text-mode login, log in there, and execute the `sudo reboot` command to reboot your system so that you can select the "(Recovery mode)" entry for your kernel. After booting using this kernel, change the `selinux=1` entry for your default kernel in the file `/boot/grub/menu.lst` to `selinux=0` to disable SELinux.

If the information in this section sounds a bit cynical, I apologize. I've used SELinux on other systems in the past with great success and minor irritation, but those options are switched on Ubuntu and Kubuntu 8.10 at the time that this book was written.

SELinux Tools

SELinux is an excellent system but is completely immature on Ubuntu and Kubuntu systems. SELinux policy definitions were originally created and maintained using text editors and activated with command-line tools, all of which are well documented in the documentation for distributions that focus on SELinux, such as Fedora (`http://fedoraproject.org/wiki/SELinux/Commands`) and Red Hat (`www.redhat.com/docs/manuals/enterprise/RHEL-4-Manual/selinux-guide`).

In addition to the standard command-line commands for SELinux and some specialized tools that Fedora/Red Hat are developing (`system-config-selinux`), a number of other tools for creating and maintaining SELinux policies are also available, some of the best known of which are the following:

- **MITRE Polgen Policy Generation Tools** (`www.mitre.org/tech/selinux`): These are GPL tools that generate SELinux policies based on analyzing program behavior and resource requirements. Red Hat has developed Polgengui, a graphical user interface to Polgen that is provided as part of the `system-config-selinux` tool (`www.redhatmagazine.com/2007/08/21`) and further simplifies policy development and generation.
- **MITRE SLAT (SELinux Analysis Tools):** These are GPL tools that help determine whether security goals are being achieved by a specified SELinux policy configuration.
- **SEEdit SELinux Policy Editor** (`http://seedit.sourceforge.net`): This is a GPL graphical tool that creates policy rules using an internal Simplified Policy Description Language (SPDL).

- **Tresys Bedrock** (`www.tresys.com/bedrock.php`): This is a commercial graphical application for policy development and maintenance that is delivered as a plug-in for the popular, open source Eclipse integrated development environment (`www.eclipse.org`).

I'm sure that other graphical tools for SELinux configuration, policy creation, and policy management are under development. These are simply the ones that I'm aware of and have experimented with. I'd appreciate hearing about other SELinux tools that you've found useful (especially once SELinux works correctly on Ubuntu and Kubuntu systems).

Additional Information about SELinux

For additional information about SELinux, see:

- **An Ubuntu hardened mailing list:** This has been set up for security development (`http://lists.ubuntu.com/mailman/listinfo/ubuntu-hardened`). See this list for the latest information about SELinux on Ubuntu and Kubuntu systems.
- **Tresys Technologies** (`www.tresys.com/education.php`): This focuses on SELinux implementation, tools, and education, Tresys has done much of the work in making SELinux available for Ubuntu and Kubuntu systems in any form.
- **National Security Agency SELinux** (`www.nsa.gov/selinux`): The originators of SELinux provide some papers, technical reports, and conference presentations on SELinux at `www.nsa.gov/selinux/info/docs.cfm`.
- **SELinux at Wikipedia** (`http://en.wikipedia.org/wiki/SELinux`): This provides information about SELinux from the consummate online encyclopedia.
- ***SELinux: NSA's Open Source Security Enhanced Linux*:** This is a book by Bill McCarty, O'Reilly, 2004, ISBN: 978-0596007164.
- ***SELinux by Example: Using Security Enhanced Linux*:** This is a book by Frank Mayer, Karl MacMillan, and David Caplan, Prentice-Hall, 2006, ISBN: 978-0131963696.
- ***Security Enhanced Linux Symposium—SELinux 2007*:** This is a book by Frank Mayer (ed), DEStech Publications, Inc., 2007, ISBN: 978-1932078671.

Summary

My rule of thumb for selecting between AppArmor and SELinux is the following: "If you are a completely anal security agency with an infinite amount of time and hordes of dedicated system policy managers, use SELinux. Otherwise, use AppArmor." Your mileage may vary. On Ubuntu and Kubuntu 8.10 (Intrepid) systems at the time that this book was written, the rule of thumb is even easier: "Don't use SELinux." It's a great system, but it is not ready for prime time on those systems at the moment.

Bonus Chapter 6 discusses various mechanisms for backing up and restoring files, starting with commands that enable you to create archives of specific files and directories, and then discussing mechanisms for backing up and synchronizing files and directories to other systems using the network. It also provides a detailed discussion of BackupPC, a powerful network-oriented application that makes it easy to automatically back up systems over the network to a central backup server.

Bonus Chapter 6

Backing Up and Restoring Files

Backups are spare copies of the files and directories that are found on a computer system, written to and stored on removable media that is preferably stored somewhere other than beside your computer. Doing backups is a time-consuming, but absolutely mandatory, task if you place any value at all on the files, e-mail, and other data that you have stored on your computer.

Backups are exactly like auto insurance policies—well, almost exactly, because they're not a legal requirement to own a computer system (in most states). You rarely need them, and you hope that you never do. They are usually just time-consuming and expensive (your time has some value, right?). However, one rainy night when you discover that you've just accidentally deleted your home directory or when a user comes to you and says that they've accidentally deleted your company's personnel records, payroll data, or the source code for your company's products, you'd better have a good answer. The right answer, of course, is, "I'll restore that from backups immediately."

It's hard to think of anything that so thoroughly combines the mundane and mandatory as backing up your data. It's boring. It's time-consuming. And, of course, it's critical. This chapter is oriented toward you as a systems administrator, regardless of how many systems you're responsible for. As system administrators, our responsibility is to provide secure, well-maintained, and rigorously backed-up systems for the benefit of the users of the computer systems we're responsible for. You should feel even more responsible if you're only supporting a user community of one (yourself), because you won't even have anyone else to blame if a catastrophe occurs. Even if you're a community of one, I'm sure that you feel that whatever you do on your computer system is important. Backups keep it safe.

This chapter explains a variety of solutions for creating backups on Ubuntu Linux systems, ranging from command-line solutions to some impressive graphical tools. It also covers the flip side of making backups, restoring files from them, which is what makes backups worthwhile in the first place.

IN THIS CHAPTER

Backups 101

Doing local backups to removable media

Using rsync for remote backups

Using backuppc for site-wide backups

Linux time machines

Backups 101

Before discussing the different tools used to create backups, it's useful to review some of the basic issues and approaches in backing up any kind of computer system. Although you may already be totally familiar with these concepts and occasionally mumble backup and restore commands in your sleep, providing a clear picture of what you're trying to accomplish in doing backups and how backup systems are usually designed provides a firm foundation for discussing the various tools discussed later in this chapter. Because I have no way of knowing whether you're reading this book because you're using Ubuntu or Kubuntu on your home computer or because you're deploying Ubuntu or Kubuntu throughout your business, I'm going to err on the side of caution and discuss many topics that are overkill for a home computing environment but are mandatory in multisystem business or academic environments.

Why Do Backups?

In an ideal world, backups would not be necessary. Computer hardware and software would always work correctly, users would never make mistakes, and I'd be a rock star. Unfortunately, in the real world, things are different. Computer system administrators and other members of an MIS/IT department do backups for many reasons, helping protect you against the following types of problems:

- Natural disasters such as fires, floods, and earthquakes that destroy computer systems.
- Hardware failures in disk drives or other storage media that make it impossible to access the data that they contain.
- System software problems such as filesystem corruption that might cause files and directories to be deleted during filesystem consistency checks.
- Software failures such as programs that crash and corrupt or delete the files that you're working on.
- Pilot error, aka the accidental deletion of important files and directories.

NOTE **Many people tend to confuse RAID (Redundant Array of Inexpensive Disks) arrays with backups. They are not the same thing at all. RAID arrays can be a valuable asset in keeping your existing data online and available in the face of disk failures, but they do not protect against any of the problems identified in the previous list. All of the drives in a RAID array will burn evenly in case of a fire or other natural disaster.**

In addition to protecting you against these sorts of problems accessing the data they, you, and any other users of your systems require, there are a variety of procedural and business reasons to back up the data on your computer systems. Complete and accurate backups provide:

- A complete historical record of your personal, corporate, or organizational business and financial data. Sadly enough, this includes serving as a source of information that you, your company, or your organization may someday need to defend itself or to prove its case in a lawsuit or other legal proceedings.
- A source of historical information about research projects and software development
- A way of preserving data that you do not need to make continuously available online but that you may need to refer to someday. This includes things like projects that you've completed, the home directories of users who are no longer using your systems, and so on.

A Few Words about Backup Media

Backups take a significant amount of time and require a significant investment in both media and backup devices. Nowadays, even home computer systems store tens or hundreds of gigabytes of information, which means that you either need to have fast, high-capacity backup devices, or you must prepare yourself for a laborious day or two of loading CDs, DVDs, or tapes. Other, more historical solutions such as Zip disks, Jazz disks, LS-120 disks, and so on provide such a small amount of storage that they're really only useful for backing up individual files, directories, or sets of system configuration files, and are therefore not discussed in this chapter.

If the mention of backup tapes causes flashbacks to mainframe computer days or old sci-fi movies, you may want to rethink that. Even for home use, today's backup tape drives are fast, store large amounts of data, are relatively cheap, and use tapes that fit in the palm of your hand. Though disk-to-disk backups are becoming more and more common, especially in networked environments, backup tapes are still quite popular and cost-efficient.

CD-Rs and DVD-Rs are eminently suitable for backups of home computer systems because they are inexpensive and typically provide enough storage for backing up the selected files and directories that compose most home backups. For home use, I prefer CD-R and DVD-R media over their rewritable brethren because of the cost difference and the fact that rewritable CDs and DVDs are only good for a limited number of writes. On the other hand, CD-Rs and DVD-R's are rarely appropriate for enterprise backups because even DVD-Rs are not large enough to back up complete systems, it's tricky to split backups across DVD-R media, and DVD-R's are relatively slow to write to. They can be useful when restoring a system because of their portability, because you can take them directly to the system you're restoring without having to move a tape drive, do a network restore, and so on. However, I personally prefer removable hard drives or tapes in enterprise or academic environments.

A final issue where backups are concerned is the need for off-site storage of all or specific sets of your backups. The history of personal, business, and academic computing is littered with horror stories about people who did backups religiously but stored them in a box beside the computer. After a file or natural disaster, all that the administrators of those systems were left with were poor excuses and unemployment benefits. Off-site storage is critical to your ability to recover from a true physical catastrophe, but it also raises another issue—the need for appropriate security in the storage location you select.

For the same reasons that you wouldn't leave the door to your house propped open and then go on vacation and wouldn't put a system that didn't use passwords on the Internet, you shouldn't store your backups in an insecure location. This is especially important if you are in charge of computer systems that are being used for business. Wherever you store your company's current and historical backup media should have a level of security comparable to wherever your computers are in the first place. Although your local cat burglar might not actively target a stack of CDs, removable disks, or storage locker full of backup tapes, any competitors you have would probably be ecstatic to be able to read and analyze the complete contents of your company's computer systems. Why not just save everybody time and mail them your source code and customer lists?

Different Types of Backups

Now that I've discussed why to do backups and some of the basic issues related to storing them, let's review the strategy behind actually doing backups. As mentioned previously, backups take time and have associated costs such as backup media, but there are a variety of ways to manage and minimize those costs.

There are three basic types of backups:

- Archive backups, which provide a complete snapshot of the contents of a filesystem at a given time
- Incremental backups, which reflect the changes to the contents of a filesystem since a previous backup
- Spot backups, which provide a snapshot of specific files or the contents of one or more important directories at a given time

NOTE **Spot backups are the most common type of backups done by home computer users, because writing a copy of your current projects, mail folders, or even your entire home directory to a few CD-Rs or DVD-Rs is relatively fast and cheap. There isn't all that much to say about this approach, because it can easily be done using drag and drop, so the rest of this section focuses on the classic backup models of archives and incremental backups. I'll discuss some techniques for doing spot backups later in this chapter.**

Archive backups, often referred to as archives or full backups, are the ultimate source for restoring data, because they usually contain a copy of every file and directory on a specific filesystem or under a certain directory on your computer at the time that the backup was done. In an ideal world, it would be great to be able to do daily archive backups simply because this would guarantee that no one could ever lose more than a day's work, regardless of the type of calamity that occurred to your computer system. Unfortunately, archive backups have some drawbacks:

- They take the maximum amount of time that backups could require because they make a copy of every file and directory on every filesystem on all of your computer systems.
- The volume of data that is preserved by an archive backup means that they use the maximum amount of space on your backup media.
- Producing the largest possible volume of backup media maximizes the amount of storage space required to store it and makes your record keeping as complex (and as critical) as it possibly could be.
- Archives are best done when no one is working on a computer system. This reduces the amount of time that it takes to do the backups (because they're not competing with anyone for computer time) and also guarantees the consistency of the files and directories that are being copied to your backup media, because nothing can be changing. This may not be a big point in a home computing environment, but in a business environment, making sure that no one is using a computer system so that you can do an archive backup is often impractical (as on systems that run 24×7services such as Web servers, database systems, and so on) or, best case, reduces the availability of a computer system to the company and your customers.

Although the advantages of archive backups as a complete record of everything are significant, these kinds of issues keep archives from being a reasonable approach to daily backups for any home computer, business, or organization. You could always do them less often than daily, but reducing the frequency of your backups increases your exposure to losing a significant amount of data if your disks fail or your computer bursts into flames.

Enter incremental backups. As mentioned before, incremental backups are backups that contain a copy of all of the files and directories that have changed on a computer system since some previous backup was done. If a problem occurs and you need to restore files and directories from backups, you can restore an accurate picture of those files and directories by first restoring from an archive backup, followed by restoring from some number of incremental backups up through your most recent ones, which should restore whatever you've backed up to the date of your most recent incremental backups. When combined with archives, incremental backups provide the following advantages:

- They help minimize the amount of disk space or other backup media required to do backups. Archives usually require large quantities of most types of backup media, while incrementals inherently require less because they aren't preserving as much data.
- They can be done more quickly, because they are copying less data than an archive backup would.
- The backup media to which incremental backups are written requires less storage space than archive backups, because there's less of it.
- Even in business and academic environments, incremental backups can be done while the computer systems and filesystems you're backing up are available for use.

Another nice feature of incremental backups is that they record changes to the files and directories on your computer systems since some previous backups, which are not necessarily archives. In corporate environments, most systems administrators organize their backup media and associated procedures in a way similar to the following:

- Archives are done infrequently, perhaps every six months or so, or just before any point at which major changes to your filesystems or computer systems are being made.
- Monthly incremental backups are made of all changes since the previous archive. If your budget and backup storage capabilities are sufficient, you usually keep the monthly incremental backups around until you do another archive backup, at which point you can reuse them.
- Weekly incremental backups are made of all changes since the previous monthly backup. You can reuse these each month, after you do the new monthly backups.
- Daily backups are made of all changes since the previous weekly backup. You can reuse these each week, after you do the new weekly backups. Some installations even just do dailies since a previous daily or the daily done on some previous day of the week.

No backup system can make it possible to restore any version of any file on a computer system. Even if you were lucky or compulsive enough to be doing daily archives of all of your computer systems, files that exist for less than a day can't be restored, and it isn't possible to restore a version of a file that is less than a day old. Sorry. When designing a backup schedule and the relationships between archive and various incremental backups, you have to decide the granularity with which you might need to restore lost files. For example, the general schedule of archives, monthlies, weeklies, and dailies doesn't guarantee that you can restore a version of a file that is newer than the previous archive. For example:

- If the file was deleted one day before the first set of monthly backups were done based on the archive, it would be present on the archive and on the weekly backups for a maximum of one month. At that point, the weekly tape containing that file would be overwritten and the newest version of the file that could be restored was the version from the archive.
- If the file was deleted one day after the first set of monthly backups were done based on the archive, it would be present on the archive and on the first monthly backup for a maximum of

seven months—a new archive would be done at that point, and the monthly tape wouldn't be overwritten until one month after the new archive. At that point, the monthly tape containing that file would be overwritten and the newest version of the file that could be restored was the version from the most recent archive.

Selecting a backup strategy is essentially a calculation of how long it will take someone to notice the absence of one or more files and request a restore, taking into account the level of service that you need to provide and the cost of various levels of service in terms of media, backup time, and storage/management overhead. Sometimes you will notice missing files immediately, such as when you accidentally delete the great American novel that you're actively working on. Other problems, such as lost files because of gradual disk failures or filesystem corruption, may not surface for a while.

Almost all backup systems generally provide automatic support for doing incremental backups since a previous incremental or archive backup. The Linux dump program, which I'll discuss in the next section, assigns different numbers to different backup "levels," and keeps track of which levels of backups have been done based on the name of the device on which the filesystem is located.

A final issue to consider when doing backups and restoring files is when to do them, and what privileges are required. It's generally fastest to do backups during off-peak hours when system usage is generally at a minimum, so that the backups can complete as quickly as possible, and when people are less likely to be modifying the files that you're backing up. In an enterprise environment, this may mean that you'll want to have a graveyard shift of operators. In this case, you'll need to think about how to make sure that operators have the right set of privileges. Being able to back up user files that may be heavily protected, or using a backup system that accesses the filesystem at the filesystem level generally requires root privileges. Many people use programs such as `sudo` (which is already our friend on Ubuntu systems) or set s-bits on privileged binaries such as backup and restore programs so that they don't have to give the administrative password to the operators or part-time staff that generally do backups at off-peak hours. See the section of Bonus Chapter 4 entitled "Customizing the `sudo` Command" for some suggestions on how to assign different privileges to different groups of individuals on your Ubuntu or Kubuntu system.

Verifying and Testing Backups

Just doing backups isn't a guarantee that you're safe from problems, unless you're also sure that the backups you're making are readable and that files can easily be restored from them. Although it's less common today, there's always the chance that the heads in a tape drive may be out of alignment. This either means that you can only read the tapes back in on the same tape drive that you wrote them on, or that they can't be read at all. You should always verify that you can read and restore files from backups using another device than the one on which they were made. You don't have to check every tape every day, but random spot checks are important for peace of mind and for job security. Similarly, tapes can just stretch or wear out from use—be prepared to replace the media used to do various types of incremental backups after some set amount of time. Nobody appreciates WORN backup media—write once, read never—even though its storage capacity is apparently infinite.

One of the problems inherent to backups is that every type of computer media has a shelf life of some period of time, depending on the type of media, the environment in which it is stored, and how lucky you are. No backup media has infinite shelf life. For example, backup tapes can last for years, but they can also be unreadable after a much shorter period of time. Long-lived media such as write-once CD-Rs and DVD-Rs are attractive because of their supposed longevity, but they have other problems, as mentioned earlier in the sidebar entitled "A Few Words about Backup Media." Media such as these may only be suited for certain types of backups, depending on whether your backup software writes to the backup

device as a filesystem or as a raw storage device. Also, no one yet knows exactly how long those types of media will last, but they certainly take up less room than almost any kind of tape or stack of hard drives. In addition to spot-checking the backup media that you are currently using, you should always make a point to spot-check old archives every few years to make sure that they're still useful.

TIP **Aside from the fact that backups can be subject to the vagaries of the device on which they're written, having those devices available when you need to restore backups is an important point to consider. It's a well-known nerd fact that many government and military sites have huge collections of backup data written on devices that don't exist anymore, such as super-low-speed tape drives and 1" or 7-track tapes. Even if the devices exist, the data is often not recoverable, because it's written in some ancient, twisted backup format, word size, and so on. When you retire a computer system, deciding if you'll ever need to restore any of its archive data is an easily overlooked issue. If you're lucky, you'll be able to read in the old archives on your new system and write them back out to some newer backup media by using some newer backup format. If you're not, you've just acquired a huge number of large, awkward paperweights that will remind you of this issue forever.**

Deciding What to Back Up

Aside from cost-saving issues like using higher-density media such as CD-ROMs for archive purposes, another way to reduce the number of old backups that you have to keep around, as well as minimize the time it takes to do them, is to treat different filesystems differently when you're backing them up. For example, system software changes very infrequently, so you may only want to back up the partitions holding your operating system when you do an archive. Similarly, even locally developed application software changes relatively infrequently, so you may only want to back that up weekly. I can count on one hand, with one finger, the number of times that I've needed to restore an old version of an application. On the other hand, you may not be so lucky.

Keeping backups of your operating system and its default applications is important, and is certainly critical to restoring or rebuilding an entire system should you ever need to do so (which is known in backup circles as a *bare-metal restore*). In terms of backups (and thanks to the excellence of the package management software available for Ubuntu and Kubuntu systems), you can usually just preserve your original installation media (or even re-retrieve it over the net) if it is ever necessary to completely restore the system software for your Ubuntu or Kubuntu system. However, if your systems run a custom kernel or use special loadable kernel modules, you should always make sure that you have a backup of your current configuration and all of the configuration files in directories such as `/etc` that describe the state of your system. You'll be glad you did if the disk on which your finely tuned and heavily tweaked version of an operating system bursts into flames late one night.

The issues in the first few sections of this chapter often give system administrators and system managers migraines. Losing critical data is just as painful if you're only supporting yourself. Thinking about, designing, and implementing reasonable backup policies, schedules, and disaster recovery plans is an important task no matter how many people will be affected by a problem. Backups are like insurance policies—you hope that you never need to use them, but if you do, they had better be available.

Backup Software for Linux

Many backup utilities are available for Ubuntu and Kubuntu systems. Most of these are traditional command-line utilities that can either create archive files or write to your backup media of choice in various formats, but some interesting open source graphical solutions are also beginning to appear.

The next few sections discuss the most common open source utilities that are used to do backups on Linux systems, grouping them into sections based on whether they create local backup files or are inherently network-aware. As discussed in the previous section, off-site storage of backups is an important requirement of a good backup strategy. In today's networked environments, off-site storage can be achieved in two basic ways: either by writing to local backup media and then physically transporting that media to another location, or by using a network-aware backup mechanism to store backups on systems that are physically located elsewhere.

Commercial Backup Solutions for Linux

Commercial application vendors have also been moving into the Linux backup space, because of the popularity of Linux in server environments. In my opinion, commercial backup solutions for Linux are unnecessary in the home computing environment, but are well worth considering if you are using Ubuntu or Kubuntu in a multisystem enterprise or academic computing environment. Before you scoff at using commercial software for backups on a free, open source operating system, consider just how critical backups are to the computers that you maintain. This book doesn't cover commercial backup solutions, but you may want to investigate the following vendors if you need or want a commercial, supported backup solution:

- **Arkeia Software:** Arkeia's Network Backup and Smart Backup applications are well-known cross-platform packages that also support backing up Microsoft Windows systems, and have been bundled with a variety of commercial Linux distributions such as Red Hat, SUSE, and Mandrake. For more information, see `www.knox-software.com`.
- **Lone Star Software:** Lone Star's LONE-TAR software has been providing backup solutions for UNIX and Linux systems for over a decade. They also offer an X Window system interface, LTX, for LONE-TAR. For more information, see `www.cactus.com`.
- **MicroLITE:** MicroLITE's BackupEDGE and RecoverEDGE software is a nice high-performance package that supports local and remote backups and provides some interesting bare-metal restore capabilities. For more information, see `www.microlite.com`.
- **Storix:** Storix's System Backup Administrator is a graphical tool for Linux (and AIX) backups that also supports Microsoft Windows backups via SMB network access and features some interesting bare-metal restore capabilities for Linux systems. For more information, see `www.storix.com`.
- **Symantec/Veritas:** Veritas's Backup Exec software, now owned by Symantec, is a well-known backup solution for Microsoft Windows systems that is also supposedly available for Linux. For more information, see `www.symantec.com/enterprise/veritas`.
- **TOLIS Group:** The BRU Backup and Restore suite is a well-known, and well-respected client/server backup solution that provides backup solutions across Linux, UNIX, Mac OS X, and Microsoft Windows platforms, and has recently ventured into the virtualization space through a partnership with Parallels. For more information, see `www.tolisgroup.com`.

None of these vendors explicitly states that it supports Ubuntu or Kubuntu, but they all advertise Linux support and I'm sure that they'd at least be happy to discuss Ubuntu or Kubuntu support with you.

Local Backup and Restore Software for Linux

The roots of the core set of Linux utilities lie in UNIX, so it's not surprising that versions of all of the classic UNIX backup utilities are available with all Linux distributions. Some of them are starting to show their age, but these utilities have been used for years and guarantee the portability of your backups from any Linux system to any other.

The classic Linux/UNIX backup utilities available in the Ubuntu and Kubuntu distributions are the following, in alphabetical order:

- **cpio:** The `cpio` utility (copy input to output) was designed for doing backups, taking a list of the files to be archived from standard input and writing the archive to standard output or to a backup device using shell redirection. The `cpio` utility can be used with filesystems of any type, because it works at the filesystem level and, therefore, has no built-in understanding of filesystem data structures.
- **dd:** The original UNIX backup utility is called `dd`, which stands for *dump device*, and it does exactly that, reading data from one device and writing it to another. The `dd` utility doesn't know anything about filesystems, dump levels, or previous runs of the program—it's simply reading data from one source and writing to another, although you can manipulate the data in between the two to do popular party tricks like converting ASCII to EBCDIC. The `dd` utility copies the complete contents of a device, such as a disk partition to a tape drive, for backup purposes. It wasn't really designed to do backups, although there are situations in which dd is the perfect tool: For example, `dd` is the tool for you if you want to copy one partition to another when a disk is failing, make on-disk copies of the partitions on a standard boot disk for easy cloning, or use an application that reads and writes directly to and from raw disk partitions that you can only backup and restore as all or nothing. Because `dd` reads directly from devices and, therefore, doesn't recognize the concept of a filesystem, individual file restores are impossible from a partition archive created with `dd` without restoring the entire partition and selecting the files that you want.
- **dump/restore:** The dump and restore utilities were designed as a pair of utilities for backup purposes and have existed for UNIX since version 6. Although `cpio` and `tar` combine the ability to write archives with the ability to extract files and directories from them and `dd` can't extract anything except an entire backup, the dump program only creates backups and the `restore` program only extracts files and directories from them. Both dump and restore work at the filesystem data structure level and, therefore, can only be used to backup and restore ext2 and ext3 filesystems (at the moment, at least). However, the dump/restore programs can accurately back up and restore any type of file that is found in ext2 and ext3 filesystems, including device-special files and sparse files (without exploding their contents and removing their "sparseness"). The dump/restore utilities can only be used to back up entire filesystems, although they have built-in support for doing incremental backups, keeping a record of which filesystems have been backed up and which level of backups has been performed for those filesystems. All of this information is tracked in an easily understood text file named `/etc/dumpdates`. Archives created with the dump utility can automatically span multiple tapes or other media if the devices support end-of-media detection, but can also span cartridge or magnetic tape media by using command-line options that tell dump the length or capacity of the tape. The most entertaining feature of the `restore` program is the ability to execute it in interactive mode, in which case it reads the information from the tape necessary to create a virtual directory hierarchy for the archived filesystem that it contains. You can then use standard commands such as `cd` to explore the list of the files on the tape and mark specific files and directories to be restored.

NOTE **The dump/restore programs are not installed as part of a default Ubuntu or Kubuntu distribution but can easily be installed using `apt-get` (both are located in the `dump` package):**

- **tar:** Probably the most widely used and well-known UNIX backup utility, the `tar` command (tape archiver) takes a list of files and/or directories to be backed up and archives those files to an output device or to standard output. The GNU version of `tar`, once known as `gtar` to differentiate it from the version of tar that came with the UNIX operating system (back when anyone cared), is yet another amazing piece of work from the Free Software Foundation. GNU `tar` provides capabilities far and above the abilities of classic UNIX `tar`, including the built-in ability to read from compressed `tar` archives created with `gzip`, support for incremental backups, support for multivolume archives, and much more. The `tar` program is filesystem-independent and accesses files and directories without needing to know their low-level data structures. The `tar` program is far and away the most popular free archiving utility available for Linux, and is used to archive almost every free software package. The DEB and RPM archive formats actually contain `tar` files that are compressed using the `gzip` utility, and files with the `.tgz` or `.tar.gz` (also gzipped `tar` files) are commonly used to distribute most Linux source code.

The utilities discussed in this section all create local archive files or write their archives to local storage devices. Of course, when you're using a network-aware operating system such as Ubuntu or Kubuntu Linux, the term *local storage devices* actually includes anything that appears to be local to your system, which therefore includes network storage that is mounted on a directory of the system that you are using. Common examples of this are NFS-mounted directories or directories that are mounted on your Linux system via Samba. Directories that are mounted over the network enable you to integrate remote storage with local backup commands in ways such as the following:

- Back up remote directories to local archives by mounting the remote directories on your local system and including them in the backups that you do.
- Write your backup files to remote storage by creating your backup archives in remote directories that are mounted on your system as local directories.

Both of these scenarios provide ways of satisfying the basic off-site requirement of backups through the use of network-mounted directories.

NOTE **Kubuntu systems provide a graphical archiving tool called Ark, which isn't discussed here because it is specific to KDE-based systems, and also because it simply provides a graphical user interface for standard command-line archiving utilities—which I think are better to learn about since that knowledge is eminently portable to any UNIX or Linux system. However, if you're interested in Ark, it is located on the K Menu ➪ Applications ➪ Utilities menu on KDE 4–based Kubuntu systems, or on the K Menu ➪ Utilities menu on KDE 3–based Kubuntu systems.**

Network-Oriented Backup Software for Linux

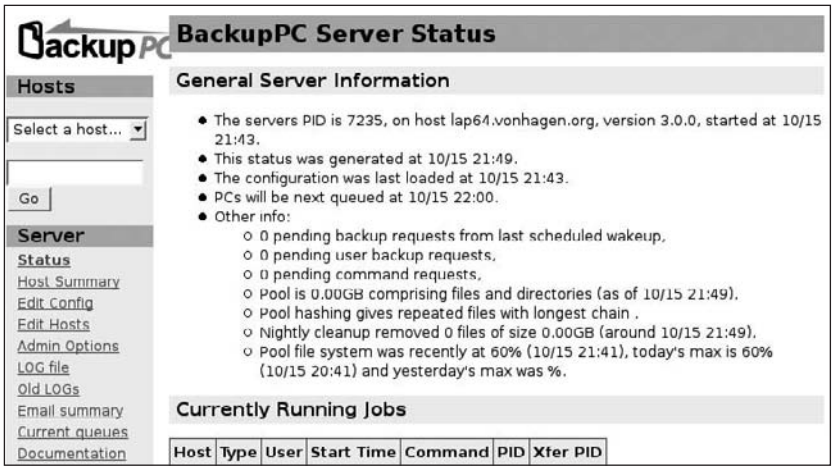
The utilities discussed in the previous section all create local archive files or write their archives to local storage devices (or storage that appears to be local). The backup utilities discussed in this section are slightly different—they are inherently network-aware and, therefore, enable you to create and manage local backups of the contents of remote systems.

The following are some of the more commonly used, network-aware backup systems that are available for Ubuntu and Kubuntu systems. There are many more, which you can find by using your package manager's search capability to search for the term *backup*. The following are my personal favorites:

- **Amanda:** The Advanced Maryland Automated Network Disk Archiver is an open source distributed backup system that was originally developed for UNIX systems at the University of Maryland in the early 1990s. Amanda makes it quite easy to back up any number of client workstations to a central backup server, supports Windows Microsoft backups via Samba, and provides a complete backup management system for your Ubuntu system. Amanda supports multiple sets of backups with distinct configurations, supports disc and tape backups, tracks backup levels and dates on its client systems, produces detailed reports that are automatically delivered via e-mail, and keeps extensive logs that make it easy to diagnose and correct the reason(s) behind most problems. Communication between Amanda clients and servers is encrypted to heighten security. Amanda is not installed by default on Ubuntu or Kubuntu systems but is available in the Ubuntu repositories and can easily be installed your favorite package management software. Amanda consists of two packages, `amanda-server` and `amanda-client`. Amanda's home Web site is at `www.amanda.org`.
- **BackupPC:** BackupPC is a nice backup system that provides a Web-based interface that enables you to back up remote systems using `smb`, `tar`, or `rsync`. Figure B6-1 shows the Web page of a sample BackupPC server. BackupPC creates backups of your remote systems that are stored and managed on your BackupPC server, and also enables authorized users to restore their own files from these archives, removing the number one source of migraines for system administrators. Configuration data for each client system is stored on the BackupPC server, which enables you to back up different types of systems using different commands or protocols, and to easily identify which remote directories or filesystems you want to back up. One especially nice feature of BackupPC is that it uses standard Linux commands on the server to create backups and, therefore, doesn't require the installation of any software on client systems, though some client-side configuration may be necessary for certain backup commands.

FIGURE B6-1

A BackupPC server's home page



- BackupPC's home page is located at `http://backuppc.sourceforge.net`. See "Installing and Using BackupPC," later in this chapter for more information about installing, setting up, and using BackupPC.

NOTE **BackupPC works fine in any modern Web browser, including FireFox and Konqueror. To emphasize this fact and toss a bone to Kubuntu fans, all of the screenshots of BackupPC's Web interface in this chapter exclude the window and browser decorations.**

- **Bacula:** Bacula is an extremely powerful set of programs that provide a scalable network backup and restore system that supports Linux, UNIX, and Microsoft Windows systems. Its power and flexibility easily match that of Amanda, but it is more flexible in terms of how and where backups are stored. Bacula is not installed by default on Ubuntu or Kubuntu systems, but it is available in the Ubuntu repositories and can easily be installed using `Synaptic`, `apt-get`, or `aptitude`. Bacula is quite powerful, but can be complex—if you're interested in exploring Bacula, you may want to start by installing the `bacula-doc` package and reading its documentation to determine if it is right for your environment. Bacula's home page is `www.bacula.org`.
- **Rsync:** Rsync (remote sync) is a command-line file and directory synchronization program that makes it easy to copy files and directories from one host to another. When both a local and remote copy of a file or directory hierarchy exist, `rsync` is able to leverage built-in features that help reduce the amount of data that needs to be transmitted to ensure that the local and remote copies of those files and directories are identical. The remote-update protocol used by the `rsync` utility enables `rsync` to transfer only the differences between two sets of files and directories. The `rsync` program is automatically installed as part of a default Ubuntu installation, but requires some configuration on the remote systems that you want to copy to your local host.

Backing Up Files to Local, Removable Media

The section of this chapter entitled "Backups 101" introduced the basic concepts of backups, many of which may seem impractical for home use. Whether or not they are really impractical depends on the problems that you want to be able to solve using your backups. If you're mostly interested in protecting yourself against disk failures or the accidental deletion of critical files that you're working on, you may not need to worry about doing archive and incremental backups—doing spot backups of important files and directories to a CD-R or DVD-R may suffice. Similarly, if you don't need to be able to restore any file from any point in time, but just need to have recent copies of your files, then spot backups of the directories that you want to back up may be sufficient, done with whatever frequency you're comfortable with. If you're not concerned about losing all of your data if your house or apartment is destroyed, then you don't have to worry about things like storing backups off-site.

The bottom line is that I can't tell you what you're comfortable with—that's up to you, and defines your backup strategy. The next few sections highlight how you can use some of the utilities mentioned earlier in this chapter (and even the standard Linux `cp` command) to create backup copies of important files.

Archiving and Restoring Files Using tar

The `tar` program is one of the oldest and most classic Linux/UNIX utilities. Though it can write to a backup device, such as a tape drive, the `tar` command is most commonly used to create archive files, such as source code, that can easily be shared with others. Archive files created using the `tar` command

typically have the .tar file extension. The GNU tar command, which is the version of `tar` found on Ubuntu and all other Linux systems, provides built-in compression capabilities, being able to automatically compress tar archives on the fly. Compressed tar archives typically have either the file extension `.tgz`, indicating that they are compressed (and can be uncompressed) using the `gzip` application, or the file extension `.tar.bz2`, indicating that they are compressed (and can be uncompressed) using the `bzip2` application. Archive files produced using the `tar` utility are typically referred to as *tarballs*.

TIP **For home use, the most popular backup method is simply dragging and dropping directories to CD-R or DVD-R media to create spot backups of those directories. The second most popular way of backing up your system is to use hard drives that you can attach to your systems via USB or FireWire ports. On the plus side, unless you're using a really small removable hard drive, this gives you a larger pool of available storage for backups than a CD or DVD, and enables you to either store more backups of important files and directories or create a single copy of each important directory on removable storage, which you can then just update each time you do backups. On the minus side, a removable hard drive is much more expensive than CD-R or DVD-R disks and is more of a pain to store off-site and retrieve each time you do backups.**

NOTE **Because of its age, you have to be kind when passing arguments to the `tar` command, because in some cases they must be specified in a particular order.**

Creating an archive file using tar is easy. For example, to create a tarball called `home_dir_backup.tgz` that contains all of the directories in `/home`, you could use commands like the following:

```
$ cd /home
$ sudo tar czvf /tmp/home_dir_backup.tgz *
```

NOTE **Note that you want to write the backup file somewhere other than the directory that you are backing up. Creating a backup file in the directory that you're working in would cause the `tar` command to back up the file that it was creating, which would both not work correctly and waste tremendous amounts of space.**

The tar options in the preceding command have the following meanings:

- `c`: Create a new archive file. If a file by the specified name already exists, it will be overwritten and its original contents will be lost.
- `z`: Compress the archive file using the same techniques used by the `gzip` application.
- `v`: Be verbose, displaying the name of every file added to the archive file as it is added.
- `f`: Write the output of the tar command to the file whose name appears as the next argument on the command line. In this example, the output of the tar command would be written to the file `/tmp/home_dir_backup.tgz`.

After a significant amount of output, the file `/tmp/home_dir_backup.tgz` will be created, containing a complete recursive copy of all files and directories under `/home`. You can then copy this file to backup media such as a CD or DVD, or to a removable hard drive.

After you've created a tarball of a given set of directories, you can easily create another tarball that only contains files and directories that have changed since a specific date (such as the date on which the first tarball was created), using commands like the following:

```
$ cd /home
$ sudo tar czvf /tmp/home_dir_backup.tgz * --newer "2006-06-23"
```

This command produces extremely verbose output, even if you drop the `v` option, which is puzzling at first. This is an artifact of the format used in tar files. Even when used with the `--newer` option, the `tar` file header must contain the complete directory structure in which it is looking for files newer than the specified date. This is necessary so that the tar command can create extracted files in the right directory location. In other words, if you use the tar command to extract the entire contents of a tarball created using the `--newer` option, it will create an empty directory hierarchy that only contains files that are newer than the date that was specified when the tarball was created.

Creating tarballs isn't much fun without being able to retrieve files from them. You can extract various things from a tarball:

- Its entire contents. For example, the following command would extract the entire contents of the tarball `home_dir_backup.tgz`, creating the necessary directory structure under the directory in which you executed the command:

  ```
  $ sudo tar zxvf home_dir_backup.tgz
  ```

- One or more directories, which recursively extracts the complete contents of those directories. For example, the following command would extract the directory `Ubuntu_Bible`, and all the subdirectories and files that it contains, from the tarball `home_dir_backup.tgz`, creating the necessary directory structure under the directory in which you executed the command:

  ```
  $ sudo tar zxvf home_dir_backup.tgz Ubuntu_Bible
  ```

- One or more specific files, which extracts only those files but creates all of the directories necessary to extract those files in their original location. For example, the following command would create the directory `Ubuntu_Bible` and extract the file `chap22.txt` from the tarball `home_dir_backup.tgz`, creating the Ubuntu_Bible directory under the directory in which you executed the command:

  ```
  $ sudo tar zxvf home_dir_backup.tgz Ubuntu_Bible/chap22.txt
  ```

For more detailed information on the `tar` command, see its online reference information (`man tar`). As one of the oldest Linux/UNIX commands, it has accumulated a huge number of command-line options over the years, many of which you will probably never use. However, command-line options are like bullets—you can never have too many.

Making an Up-to-Date Copy of a Local Directory Using cp

If you're only backing up a few directories and are primarily concerned with keeping copies of the files that you are actively working on, it's often simplest to just keep copies of those directories on removable media. The traditional Linux/UNIX `cp` command provides options that make it easy to create a copy of a specified directory and then to subsequently update only files that have been updated or that do not already exist in the copy. For example, to back up all of the directories in /home to a removable drive mounted at `/media/LACIE` (LACIE is a popular manufacturer of prepackaged USB hard drives), you could use a command like the following:

```
$ sudo cp -dpRuvx /home /media/LACIE/home
```

The `cp` options in this command have the following meanings:

- `d:` Don't dereference symbolic links, that is, copy them as symbolic links instead of copying what they point to.

- p: Preserve modes and ownership of the original files in the copies.
- R: Copy the specified directory recursively.
- u: Copy files only if the original file is newer than an existing copy or if no copy exists.
- v: Display information about each file that is copied. (You may not want to use this option, but it's interesting, at least the first few times you do this.)
- x: Don't follow mount points to other filesystems.

After running this command, you will have a copy of every directory under /home on your system in the directory /media/LACIE/home. You can then detach your removable drive and store it somewhere safe (preferably off-site). Any time that you want to update your backup, retrieve the drive and update this copy at any time by simply rerunning this command.

Making an Up-to-Date Copy of a Remote Directory Using rsync

As mentioned earlier, rsync is a commonly used command-line utility that enables you to push or pull files to or from remote systems. The rsync program must be configured on the remote systems before you can push or pull file or directories to or from those systems.

To use rsync on an Ubuntu or Kubuntu system, you must first enable it so that the system starts rsync as a background process, and then also modify the rsync configuration file to add entries for specific directories that you want to be able to read from and write to remotely. To enable rsync, edit the file /etc/defaults/rsync using your favorite text editor and a command like the following:

```
$ sudo emacs /etc/default/rsync
```

In the line that begins with RSYNC_ENABLE, change false to true, and then save the updated file. Next, create the rsync configuration file before actually starting the rsync daemon.

NOTE **Most Linux systems use an Internet service manager such as inetd or xinetd to manage incoming requests for on-demand services such as ftp, tftp, rsync, and vnc. These Internet service managers automatically start the appropriate daemon when an incoming request is received. Although these Internet service managers are available in the Ubuntu repositories, they are not installed by default. On Ubuntu and Kubuntu systems, a specific system startup file that starts rsync in daemon mode is provided as /etc/init.d/rsync. If you subsequently install xinetd and want to use it to manage rsync requests, you will want to disable this file and create the file /etc/xinetd.d/rsync to make sure that the rsync service is enabled on your system.**

The /etc/defaults/rsync file just determines whether rsync is enabled or not. The actual configuration information for rsync itself is stored in the file /etc/rsyncd.conf, which does not exist by default on an Ubuntu or Kubuntu system. To create this file, use your favorite text editor and a command like the following:

```
$ sudo emacs /etc/rsyncd.conf
```

A minimal rsync configuration file that contains a definition remotely synchronizing the directories under /home on your system would look something like the following:

```
uid = root
transfer logging = true
log format = %h %o %f %l %b
log file = /var/log/rsyncd.log
hosts allow = 192.168.6.255/3
```

```
[homes]
    path = /home
    comment = Home Directories
    auth users = wvh
    secrets file = /etc/rsyncd.secrets
```

The first section of this file sets parameters for how the `rsync` dameon runs. In order, the `rsync` daemon runs as root (`uid`), logs all transfers (`transfer logging`), uses a specific log file format (`log format`) and log file (`log file`), and allows access from any host whose IP address is on the 192.168.6.*x* subnet (`hosts allow`). The second section of this file identifies a synchronizable entity known as home that maps to the directory `/home` on that system. Synchronization to or from this directory is done as the user `wvh`, whose password must be supplied in the file `/etc/rsyncd.secrets`.

After saving this file, use the `sudo` command and your favorite text editor to create the file `/etc/rsync.secrets`, with a command like the following:

```
$ sudo emacs /etc/rsyncd.secrets
```

This file should contain an entry for each `auth users` entry in the `/etc/rsync.conf` file, in this case, `wvh`. Each entry in this file contains the name of a user, a colon, and the plain-text password for that user, as in the following example:

```
wvh:hellothere
```

Next, save this file and make sure that it is readable only by the root user on your system using a command like the following:

```
$ sudo chmod 600 /etc/rsyncd.secrets
```

You can now start the `rsync` daemon using the following command:

```
$ sudo /etc/init.d/rsync restart
```

You can now create a local copy of the `/home` directory on your system using a command like the following, where `system-name-or-addr` is the name or IP address of the system on which you just configured the `rsync` daemon:

```
$ rsync -Havz system-name-or-addr::home /media/LACIE/home
```

The arguments to the `rsync` command in this example have the following meaning:

- `H:` Preserve hard links if these exist in any directories that are being copied.
- `a:` Use archive mode, which preserves ownership, symbolic links, device files, and so on, and is essentially a shortcut that saves you specifying several other options.
- `v:` Be verbose, identifying each file that is copied or considered for copying. (You may not want to use this option, but it's interesting, at least the first few times you run `rsync`.)
- `z:` Use compression when transferring files, which improves throughput.

If you have problems using `rsync`, you should check the `/var/log/rsyncd.log` file (on the system that you are trying to retrieve files from) for error messages and hints for resolving them. If you are not using

the verbose option on the host where you are retrieving these files, you may want to use it to see if you can identify (and resolve) any other errors that the host that is trying to retrieve files is reporting.

NOTE **The `rsync` configuration file created in this section is just a minimal example, and is not particularly secure. For details about all of the options available in an `rsync` configuration file and information about making `rsync` more secure, see the man page for the `rsyncd.conf` file (`man rsyncd.conf`).**

Installing and Using BackupPC

This section explains how to install, configure, and use the BackupPC package to back up a variety of hosts on your local network to a central Ubuntu or Kubuntu server. Introduced earlier in this chapter, BackupPC is a great application that is both easy to use for a system administrator and empowering for any authorized user. Any authorized user can initiate backups of the machines that they have admin rights to and can also restore files from existing backups of those machines, all using a convenient Web interface.

If you have more than one machine on your home network, or if you're working in a multimachine enterprise or academic environment, the BackupPC software is well worth a look. Its Web-based interface is easy to set up and use; various types of supported backups are easy to configure, initiate, and monitor; it can back up your Linux, UNIX, Windows, and Mac OS X systems; and the fact that it doesn't require that you install any special software on the systems that you want to back up makes BackupPC a great package.

BackupPC supports four different backup mechanisms (known in the BackupPC documentation as *backup transports*) to enable you to back up different types of systems. These are the following:

- **rsync:** Back up and restore via `rsync` via `rsh` or `ssh`. This is a good choice for backing up Linux, UNIX, or Mac OS X systems, and you can also use it to back up Microsoft Windows systems that support `rsync`, such as those running the Cygwin Linux emulation environment.
- **rsyncd:** Back up and restore via the `rsync` daemon on the client system. This is the best choice for Linux, UNIX, and Mac OS X systems that are running an `rsync` daemon. You can also use this mechanism to back up Microsoft Windows systems that support `rsyncd`, such as those running the Cygwin Linux emulation environment.
- **smb:** Back up and restore using the `smbclient` and the SMB protocol on the `backuppc` server. This is the best (and easiest) choice to use when backing up Microsoft Windows systems using `backuppc`, and you can also use it to back up Mac OS X systems or Linux and UNIX systems that are running a Samba server.
- **tar:** Back up and restore via `tar`, `tar` over `ssh`, `rsh`, or `nfs`. This is an option for Linux, UNIX, and Mac OS X systems. You can also use this mechanism to back up Microsoft Windows systems that support `tar`, `ssh`, `rsh`, and/or `nfs`, such as those running the Cygwin Linux emulation environment.

A default backup transport value for all backups is set in the primary BackupPC configuration file, `/etc/backuppc/config.pl`. The specific mechanism used to back up any particular host can be identified in that host's configuration file, as discussed later in the sections entitled "Defining a Backup Using rsyncd" and "Defining a Backup Using SMB."

NOTE **Although `backuppc` does a great job of backing up systems running Microsoft Windows and Mac OS X, you should be aware of a few issues. First, `backuppc` is not suitable for backing up Windows systems so that you can do a bare-metal restore. `Backuppc` uses the `smbclient` application on your Ubuntu system to back up Windows disks, so it doesn't back up Windows ACLs and can't open files that are locked by a Windows client that is currently running (such as, most commonly, things like Outlook mailboxes). Similarly, `backuppc` doesn't preserve Mac OS file attributes. See `http://backuppc.sourceforge.net/faq/limitations.html` for a list of current limitations in using `backuppc`. It's a surprisingly short document!**

For additional information about BackupPC, you can join various mailing lists that are available through the SourceForge BackupPC site. These lists are the following:

- **backuppc-announce:** A low-traffic list where announcements of new versions are posted (`http://lists.sourceforge.net/lists/listinfo/backuppc-announce`)
- **backuppc-devel:** The mailing list for BackupPC developers. (`http://lists.sourceforge.net/lists/listinfo/backuppc-devel`)
- **backuppc-users:** The standard mailing list where you can ask and answer questions about using BackupPC (`http://lists.sourceforge.net/lists/listinfo/backuppc-users`)

Installing and Configuring BackupPC

Special-purpose backup solutions such as BackupPC aren't installed as part of a default Ubuntu or Kubuntu installation because they're probably overkill for most people. However, the various package management tools provided on Ubuntu and Kubuntu systems make it easy to install BackupPC (located in the `backuppc` package) and the other software packages that it requires.

Because BackupPC uses a Web-based interface, it requires that you have installed (and configured) a Web server on your system, preferably Apache2. Similarly, because BackupPC can send e-mail notifications for various backup-related events, it requires a Mail Transfer Agent (MTA). If one is not already installed on your system, installing the `backuppc` package will install `exim`, which isn't my personal choice as far as MTAs go. If you're going to install BackupPC, I'd suggest installing (and configuring) Postfix first, as described in Bonus Chapter 9, "Setting Up a Mail Server." After doing so, BackupPC will detect that an MTA is installed on your system and will no longer try to install `exim` to satisfy dependencies.

Once the `backuppc` package has been installed on your system, your package management software will run its configuration script. During this phase, the various package managers ask which Web server configuration you're using so that they can configure `backuppc` appropriately, and automatically run a script that sets up the initial account that you will use to access `backuppc` via that Web server. This process appears a dialog like the one shown in Figure B6-2, which tells you the initial password for the Web-based BackupPC interface.

This dialog is, of course, different for each package management tool—Figure B6-2 shows the screen that you would see if installing the `backuppc` package using the aptitude package management tool. See Chapter 23 for more information about the various package management utilities that are available for Ubuntu and Kubuntu systems.

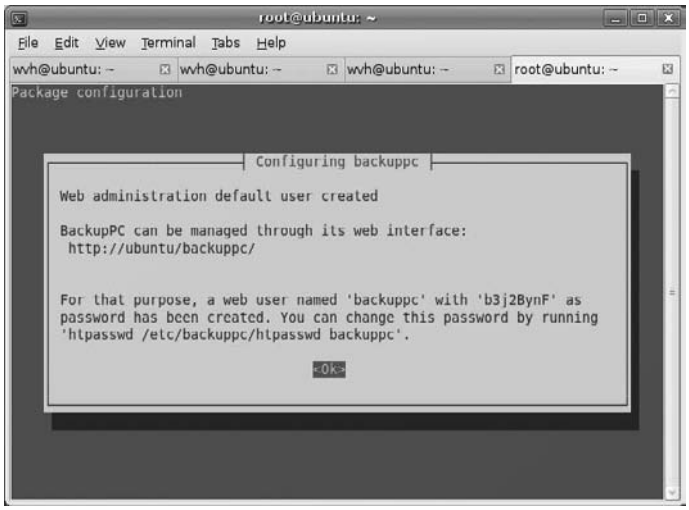
Once you see a dialog like the one shown in Figure B6-2, write down the password for the `backuppc` interface and press Return or click Forward to proceed. Unless you're good at remembering passwords like "b3j2Bynf", I'd suggest immediately changing the default password to one of your standard adminis-

trative passwords, which you can do via a command like the following (which the dialog in Figure B6-2 also explains):

```
sudo htpasswd /etc/backuppc/htpasswd backuppc
```

FIGURE B6-2

The initial password for the backuppc package



Once the remainder of the installation and configuration process completes, you're almost ready to back up the system you're using and the other systems on your network. While installing `backuppc` is supposed to integrate it with your Web server, that doesn't actually seem to be the case. To integrate `backuppc` with your Apache Web server, do the following:

1. Copy the file `/etc/backuppc/backuppc.conf` to the `directory /etc/apache2/sites-available`.
2. Create a symbolic link from the file `/etc/apache2/sites-available/backuppc.conf` to a file in the directory `/etc/apache2/sites-enabled`. I usually name this symbolic link `backuppc` so that it's clear exactly what site/application is enabled. You can do this by executing the following commands:

   ```
   cd /etc/apache2/sites-enabled
   sudo ln -s ../sites-available/backuppc.conf backuppc
   ```

3. Restart Apache using a command like the following:

   ```
   sudo /etc/init.d/apache2 reload
   ```

You should now be able to connect to the BackupPC Web interface via the `/backuppc` URL relative to your Web server as the user `backuppc` with whatever your current BackupPC password is (in other words, something like `www.vonhagen.org/backuppc`). This will display a screen like the one shown in B6-1.

On Ubuntu and Kubuntu systems, BackupPC stores its configuration information in two locations. General backuppc configuration information and passwords are stored in files in the directory /etc/backuppc. Backup files themselves and host-specific backup configuration information is stored in subdirectories of /var/lib/backuppc.

TIP

Backups of a single system take a significant amount of space, which is only compounded when you begin to back up other hosts to a central backup server. If you didn't specify using logical volumes when you installed your Ubuntu or Kubuntu system, you may want to add a new disk to your system before starting to use BackupPC and format that disk as a logical volume. You can then copy the default contents of /var/lib/backuppc to the new disk (preserving file permissions and ownership), and mount that disk on the directory /var/lib/backuppc on the system that you are using for backups. When you need more space to store backups in the future, this will enable you to add other disks to your system and add their space to the logical volume used to store backups. BackupPC also provides an archive capability that enables you to migrate old backups to other hosts for archival purposes, freeing up disk space on your primary backup server. Though not discussed in this book, setting up archive hosts is discussed in the BackupPC documentation—which is great, by the way!

Identifying Hosts to Back Up

Each host that you want to back up must be identified in the file /etc/backuppc/hosts. Like all BackupPC configuration files, this file is easy to update. Any characters in any lines in this file that follow a hash mark are comments, which help explain the meaning of the various fields used in the file. A minimal backuppc configuration file looks like the following:

```
host            dhcp    user            moreUsers
localhost       0       backuppc
```

The first noncomment line in /etc/backuppc/hosts defines the names of the various fields in each line, and should therefore not be modified. (This is the line beginning with the word "host" in the example.) All other lines represent entries for hosts that will be backed up. The first actual host entry, for localhost, is a special entry used for backing up system configuration information on the backuppc server, and should not be changed.

The fields in each entry that define a host have the following meanings:

- The first field identifies a particular machine, by hostname, IP address, or NetBios name.
- The second field should be set to 0 for any host whose name can be determined by DNS, the local hosts file, or an nmblookup broadcast. This field can be set to 1 to identify systems whose names must be discovered by probing a range of DHCP addresses, as is the case in some environments where DHCP and WINS are not fully integrated. Setting this field to 1 requires changes in the host-specific configuration file's $Conf{DHCPAddressRanges} variable to define the base IP address and range of IP addresses that should be probed.
- The third field identifies the name of the person who is primarily responsible for backing up that host. This primary user will receive e-mail about the status of any backup that is attempted. I tend to leave this as the backuppc user, so that this user maintains an e-mail record of all backup attempts, but you can set this to a specific user if you wish.
- The fourth field (which is optional) consists of one or more users who also have administrative rights to initiate backups or restore files for this machine. The names of multiple users must be separated by a comma.

As an example, the hosts file on one of my backuppc servers looks like the following:

```
host            dhcp    user            moreUsers
localhost       0       backuppc
192.168.6.64    0       backuppc        wvh
64bit           0       backuppc        wvh,djf
64x2            0       backuppc        juser
win2k           0       backuppc        wvh,djf
```

TIP **The backuppc program checks the timestamp on the /etc/backuppc/hosts files each time the backuppc process wakes up, and reloads this file automatically if the file has been updated. For this reason, you should not save changes to the hosts file until you have created the host-specific configuration files, as described in the examples in the next two sections. If the backuppc process reloads the hosts file before you have created the host-specific configuration data and another authorized user initiates a backup of this system, you will either back up the wrong thing or a backup failure will occur. You can always make changes to the hosts file and leave them commented out (by putting a # as the first character on the line) until you have completed the host-specific configuration.**

Defining a Backup Using rsyncd

The section earlier in this chapter entitled "Making an Up-to-Date Copy of a Remote Directory Using rsync" explained how to set up rsync in daemon mode on an Ubuntu or Kubuntu system and how to define synchronization entries that can be remotely accessed via rsync. The sample rsync configuration file created in that section defined a synchronization entry called homes that would enable an authorized user to synchronize the contents of all directories under /home on a sample Ubuntu system. We'll use that same configuration file in the example in this section.

The previous section showed how to define entries in the /etc/backuppc/hosts file for the various hosts that you want to back up via backuppc. The first step in host-specific configuration is to use the sudo command to create a directory to hold host-specific configuration data, logs, and so on. Throughout this section, I'll use the sample host entry 64bit, which I defined in the section entitled "Identifying Hosts to Back Up" as an example.

The first step in host-specific configuration is to use the sudo command to create the directory /var/lib/backuppc/64bit, as in the following command:

```
sudo mkdir /var/lib/backuppc/64bit
```

Next, use the sudo command and your favorite text editor to create a host-specific configuration file named config.pl in that directory, using a command like the following:

```
sudo emacs /var/lib/backuppc/64bit/config.pl
```

The contents of this file should be something like the following;

```
$Conf{XferMethod} = 'rsyncd';
$Conf{CompressLevel} = '3';
$Conf{RsyncShareName} = 'homes';
$Conf{RsyncdUserName} = 'wvh';
$Conf{RsyncdPasswd} = 'hellothere;
```

The first line identifies the backup mechanism used for this host as rsyncd, which overrides the default backup mechanism specified in the generic /etc/backuppc/config.pl file. The second line sets the

compression level for this host's backups to level 3, which provides a good tradeoff between the CPU load and time required to do compression and the amount of compression that you actually get. The last three entries in this file correspond to the synchronization entry in the sample `rsyncd.conf` and associated `rsyncd.secrets` file created in "Making an Up-to-Date Copy of a Remote Directory Using rsync" earlier in this chapter.

TIP **When using BackupPC to do automated backups, I like to create a separate authorized user to use `rsync` for backup purposes, so that the system logs show who actually requested a remote sync operation. To do this, you would add this user (I usually use `backuppc`) to the `auth users` entry in the remote host's `/etc/rsyncd.conf` file and create an appropriate username/password pair in the remote host's `/etc/rsyncd.secrets` file. You would then modify the host-specific `backuppc` configuration file to use this username and password. I didn't do this here for simplicity's sake, but doing this would provide more accurate log data on the client system.**

If the remote system uses an `rsync` binary other than the default `/usr/bin/rsync` or the `rsync` program is listening on a port other than the standard port (873), you should add correct definitions for these to the host-specific configuration file. The default settings for the associated configuration parameters are the following;

```
$Conf{RsyncdClientPort} = 873;
$Conf{RsyncClientPath} = '/usr/bin/rsync';
```

Next, change the ownership and group of the `/var/lib/backuppc/64bit` directory to `backuppc` and change the protection of the configuration file `/var/lib/backuppc/64bit/config.pl` so that it is not publicly readable (because it contains password information) using the following commands:

```
sudo chmod -Rv backuppc:backuppc /var/lib/backuppc/64bit
sudo chmod 600 /var/lib/backuppc/64bit/config.pl
```

The last step in creating a host-specific backup definition for `backuppc` is to cause the `backuppc` process to reread its configuration data, which you can do by explicitly reloading the configuration file, by explicitly restarting the `backuppc` process, or by sending the associated process a hang-up (HUP) signal. You can force `backuppc` to reload the configuration file by using the following command:

```
sudo /etc/init.d/backuppc reload
```

The definition for your backup host can now be selected via the BackupPC Web interface. At this point, you can follow the instructions in the section entitled "Starting Backups in backuppc" to back up this host.

The example in this section only backs up the home directories of users on the remote machine. To recursively back up other directories, you would simply create other synchronization entities for those directories in the remote host's `/etc/rsyncd.conf` file, and then add entries for those synchronization entities to the host-specific configuration file. For example, to back up synchronization entries named `/`, `homes`, and `/boot`, you would change the host-specific `RsyncShareName` entry to look like the following:

```
$Conf{RsyncShareName} = ['/', 'homes', '/boot'];
```

If you back up multiple filesystems or synchronization points, you may create a custom set of arguments to the `rsync` command in the host-specific configuration file. This enables you to add options such as `--one-file-system`, which causes `backuppc` to back up each filesystem separately, simplifying restores. You can also add options to exclude certain directories from the backups, which you will certainly want to do if you are backing up a remote system's root directory (`/`), as in following examples:

```
$Conf{RsyncArgs} = [
          # original arguments here
```

```
            '--one-file-system',
            '--exclude', '/dev',
            '--exclude', '/proc',
            '--exclude', '/media',
            '--exclude', '/mnt',
            '--exclude', '/lost+found',
];
```

These settings would prevent backups of `/dev`, which contains device nodes and is dynamically populated at boot time on modern Linux systems; `/proc`, which is the mount point for an in-memory filesystem that contains transient data; directories such as `/media` and `/mnt` on which removable media is often temporarily mounted; and `/lost+found`, which is a directory used during filesystem consistency checking. You can also exclude directories from `rsync` backups using the `BackupFilesExclude` directive, as in the following example:

```
$Conf{BackupFilesExclude} = ['/dev', /proc', '/media', '/mnt', '/lost+found'];
```

The `backuppc` program reads the configuration settings in `/etc/backuppc/config.pl` first, and then loads host-specific configuration settings, which enables the `/etc/backuppc/config.pl` file to provide default settings for all backups. After you have used `backuppc` for a while and are comfortable with various settings, you may want to consider modifying the default settings in the `/etc/backuppc/config.pl` file for configuration variables such as `$Conf{RsyncArgs}`, `$Conf{BackupFilesExclude}`, and `$Conf{CompressLevel}`, to minimize the number of entries that you have to create in each of your host-specific configuration files.

Defining a Backup Using SMB

The section of this chapter entitled "Identifying Hosts to Back Up" showed how to define entries in the `/etc/backuppc/hosts` file for the various hosts that you want to back up via `backuppc`. The first step in host-specific configuration is to use the `sudo` command to create a directory to hold host-specific configuration data, logs, and so on. Throughout this section, I'll use the sample host entry `win2k` from the sample hosts file as an example. As you might gather from its name, this is indeed a system running Microsoft Windows 2000. There's no escaping from the Borg.

The first step in host-specific configuration is to use the `sudo` command to create the directory `/var/lib/backuppc/win2k`, as in the following command:

```
$ sudo mkdir /var/lib/backuppc/win2k
```

Next, use the `sudo` command and your favorite text editor to create a host-specific configuration file named `config.pl` in that directory, using a command like the following:

```
$ sudo emacs /var/lib/backuppc/win2k/config.pl
```

The contents of this file should be something like the following:

```
$Conf{XferMethod} = 'smb';
$Conf{CompressLevel} = 3';
$Conf{SmbShareName} = ['wvh', "djf'];
$Conf{SmbShareUserName} = 'backuppc;
$Conf{SmbSharePasswd} = hellothere';
```

The first line identifies the backup mechanism used for this host as `smb`, which overrides the default backup mechanism specified in the generic `/etc/backuppc/config.pl` file. The second line sets the

compression level for this host's backups to level 3, which provides a good tradeoff between the CPU load and time required to do compression and the amount of compression that you actually get. The last three entries in this file define the Windows shares that you want to back up, the name of an authorized user who has access to these shares, and the password for that user.

TIP **When using `backuppc` to back up Microsoft Windows systems, you should create a Windows user that you will only use to do backups and then add this user to the standard Windows Backup Operators group. This prevents you from having to put your Windows administrator password in the `backuppc` configuration files. Even though you'll protect those files so that randoms can't read them, the fewer places where you write down a password, the better, especially one with the keys to your entire Windows kingdom.**

Next, change the ownership and group of the `/var/lib/backuppc/win2k` directory to `backuppc` and change the protection of the configuration file `/var/lib/backuppc/win2k/config.pl` so that it is not publicly readable (because it contains password information) using the following commands:

```
sudo chmod -Rv backuppc:backuppc /var/lib/backuppc/win2k
sudo chmod 600 /var/lib/backuppc/win2k/config.pl
```

The last step in creating a host-specific backup definition for `backuppc` is to cause the `backuppc` process to reread its configuration data, which you can do by explicitly reloading the configuration file, by explicitly restarting the `backuppc` process, or by sending the associated process a hang-up (HUP) signal. You can force BackupPC to reload its configuration file by using the following command:

```
sudo /etc/init.d/backuppc reload
```

The definition for your backup host can now be selected via the `backuppc` Web interface. At this point, you can follow the instructions in the section entitled "Starting Backups in backuppc" to back up this host.

The example in this section only backs up shares that correspond to the home directories of selected users on the remote machine. As mentioned earlier in this chapter, `backuppc` backups do not support bare-metal restores of Windows systems, and I therefore typically don't back up shares such as C$, which is a default Windows share that represents your system's boot drive. You may find it useful to do so to make sure that you have backup copies of drivers, the registry, and so on, but I find it simpler to start from scratch when reinstalling Windows. Windows systems accumulate so much crap in their filesystems over time that doing a fresh installation from your distribution media often frees up a surprising amount of space. If you have several identical systems, restoring partition images created with Norton Ghost or the Linux `partimage` or `g4u` utilities is always the fastest way to rebuild a Windows system without having to locate the drivers for every device that you will ever want to use with your rebuilt system and reinstalling all of your favorite applications.

TIP **The `backuppc` program reads the configuration settings in `/etc/backuppc/config.pl` first, and then loads host-specific configuration settings, which enables the `/etc/backuppc/config.pl` file to provide default settings for all backups. After you have used `backuppc` for a while and are comfortable with various settings, you may want to consider modifying the default settings in the `/etc/backuppc/config.pl` file for configuration variables, such as `$Conf{CompressLevel}`, to minimize the number of entries that you have to create in each of your host-specific configuration files.**

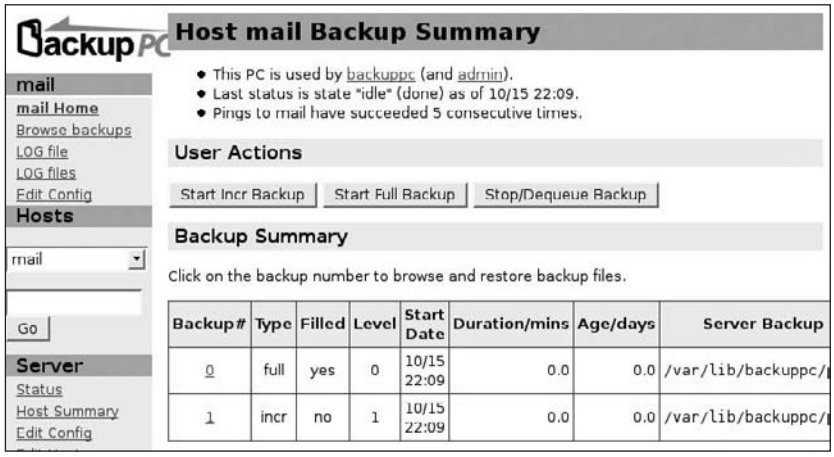
Starting Backups in BackupPC

Thanks to BackupPC's Web orientation, starting backups, viewing the status of those backups, and checking the backup history for any host are impressively easy. To start a backup in BackupPC, connect to the BackupPC Web interface, using the URL `http://hostname/backuppc`, where `hostname` is the name

of the host on which the `backuppc` server is running. A dialog appears in which you are prompted for the login and password of an authorized user. Once you enter the `user/password` combination for a user listed in the file `/etc/backuppc/htpasswd`, the BackupPC server's home page appears, as shown long ago in Figure B6-1. Once this screen appears, click the Select a host drop-down box and select one of the hosts from the list that appears. Selecting the name of any host takes you to a summary page for that host, which provides status information, lists authorized users who can back up and restore files to this host using BackupPC, and enables you to browse existing backups for that host in order to locate and restore files, as shown in Figure B6-3.

FIGURE B6-3

A per-host page in the BackupPC Web interface



NOTE **Each system's home page displays the subject of the last e-mail sent to the owner of this host. E-mail is only sent occasionally, so seeing a historical problem report does not mean that the problem is still occurring.**

Once this page appears, you can scroll down on the page to see additional status information about available backups, any transfer errors that occurred during backups, and other tables that show the status of the pool where backup files are archived and the extent to which existing backups have been compressed to save disk space.

To start a backup, click either Start Full Backup to start a full (archive) backup of the system, or Start Incr Backup to start an incremental backup containing files that have changed since the last full backup. A confirmation page appears.

Clicking Start Full Backup (or Start Incr Backup for an incremental backup) queues the backup and displays a link that you can click to return to the main page for that host to monitor the state of the backup.

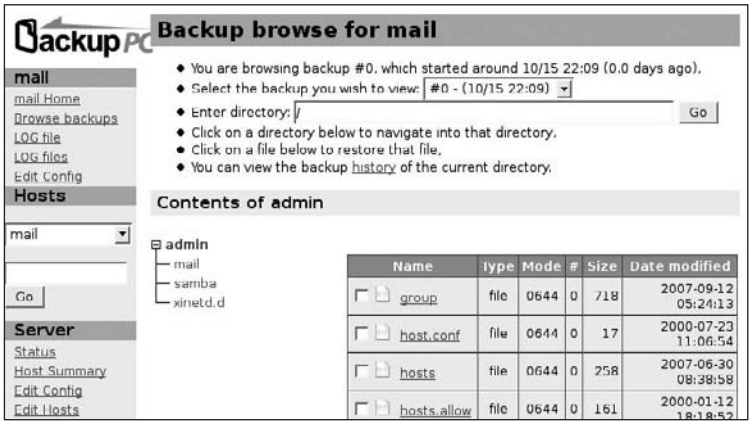
Restoring from Backups in BackupPC

Thanks to BackupPC's Web orientation and the fact that its backups are stored online on the backup server, restoring files from BackupPC backups can be done online, by any authorized user whose name is associated with that host in the `/etc/backuppc/hosts` file. BackupPC enables you to browse through online backups, interactively select the files and directories that you want to restore, and restore them in various ways.

To begin restoring files or directories, click the name of the full or incremental backup in which they are located. A screen like the one shown in Figure B6-4 appears.

FIGURE B6-4

Restoring files and directories in the backuppc Web interface



The bottom of the screen shown in Figure B6-4 displays a hierarchical listing of the files and directories that are contained in the full or incremental backup that you selected. If you selected an incremental backup, the contents of that incremental backup are overlaid on the contents of the previous full backup to give you an accurate snapshot of the contents of your system when the backup was done. You can drill down into the backup by selecting directories from the tree view at the left, or you can drill down into individual directories by selecting from the view of the current directory shown at the right of the main window.

To mark files and directories for restoration, select the checkbox beside their names. Once you have selected all of the files and directories that you want to restore, scroll to the bottom of the restore page and click restore selected files. A page that enables you to specify how you want to restore those files appears, as shown in Figure B6-5.

You have three options when restoring files using the BackupPC Web interface:

- **Direct restore:** Selecting this option restores files directly to the host from which they were backed up. When doing a direct restore, you have the option of restoring files in the locations from which they were originally backed up, or into a subdirectory that BackupPC will create for you if it does not already exist. (The latter is almost always a good idea so that you don't accidentally overwrite any files that you don't actually mean to.) To select this option, enter the name of any subdirectory that you want to use (I usually specify one called `tmp`) and click Start restore.

- **Download Zip archive:** Selecting this option restores the selected files and directories into a `.zip`-format archive that you can download to your desktop and manually extract the contents of. When selecting this option, you can optionally specify the compression level used in the `.zip` file, which can be important if you are restoring large numbers of files. To select this option, click Download Zip file.
- **Download Tar archive:** Selecting this option restores the selected files and directories into a `tar`-format archive that you can download to your desktop and manually extract the contents of. To select this option, click Download Tar file.

FIGURE B6-5

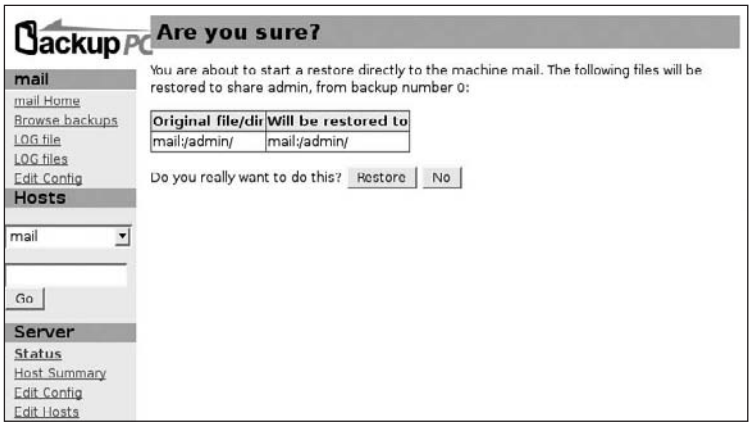
Specifying the restore method in BackupPC's Web interface



If you selected the Direct restore option, `backuppc` displays a confirmation screen like the one shown in Figure B6-6.

FIGURE B6-6

Confirming a restore directly to a remote system



This lists the files and directories that you selected for restoration and confirms the location to which they will be restored, including the name of any subdirectory that you specified. To proceed, click Restore.

If you selected the Zip or Tar archive options, the BackupPC application displays your Web browser's standard file download dialog after the archive file has been created.

As you can see from this section (and the preceding sections), BackupPC provides a powerful, flexible interface for backing up and restoring files on many different systems to a single `backuppc` server. All you need are a few configuration files and sufficient disk space, and lost files (and the lost time that is usually associated with them) can be a thing of the past.

Linux Time Machines

Apple's Mac OS X is a great computing environment—it provides an easy-to-use graphical interface on top of a powerful, UNIX-like operating system. Remind you of anything? Apple has always provided innovative, user-friendly software, and with the release of Mac OS X 10.5, aka Leopard, they introduced an impressively friendly and powerful backup package known as Time Machine. Time Machine automatically tracks changes to the files and directories on your system and backs up multiple snapshots of your system to an external (or networked) storage device. Should you accidentally delete a file or directory or simply want to consult an older version, you can explore those snapshots graphically and examine or restore a version of that data from any point in Time Machine's archives.

As the saying goes, "Imitation is the sincerest form of flattery" (Charles Colton, 1820). Not to be outdone by mere commercial software, Linux-based packages that provide equivalent functionality have begin to appear, most of which are graphical wrappers over the `rsync` utility (discussed earlier in this chapter), with various bells and whistles to make it easy to browse "snapshots," identify available versions of specific files, extracts specific versions of those files, and so on. These sorts of utilities put an impressively useful face on traditional backups, and Apple deserves major kudos from conceptualizing and automating backups to the point where they are actually user-friendly. Like Mac OS X's Time Machine software,

similar Linux backup mechanisms are best used with a networked resource or, in a single-system environment, an external drive. These backups can't protect against catastrophic problems, such as a fire or similar tragedy, unless you also "back up the backups" to an off-site location. If I had a nickel for every time that someone lost both their computers systems and their backups in a fire because they were stored side by side, I'd have a bunch of singed, smoky nickels. Don't add one to my collection.

The following are well worth a look to do snapshot-style backups on Ubuntu or Kubuntu systems:

- **Back In Time**: Back In Time is a really nice GTK-based package that provides a nice, date-oriented snapshot browser and a single-button restore mechanism. Figure B6-7 shows the Back In Time snapshot browsing and restore interface. Back In Time is my second favorite of the Linux Time Machine reimplementations, and is available as a Ubuntu/Kubuntu DEB package through its home page at `www.le-web.org/back-in-time`.
- **FlyBack:** Originally a Google Summer of Code project, FlyBack is the least active of the Linux Time Machine–like applications that I'm aware of, but works fine. FlyBack is installed from Subversion or a tarball, which is not exactly "grandfather-friendly," but it's worth a look and offers some handy features in its configuration interface, such as the ability to automatically identify networked filesystems as possible snapshot storage locations. Figure B6-8 shows the FlyBack snapshot creation and restore interface. For more information about FlyBack, see `http://code.google.com/p/flyback` and `http://flyback-project.org`.

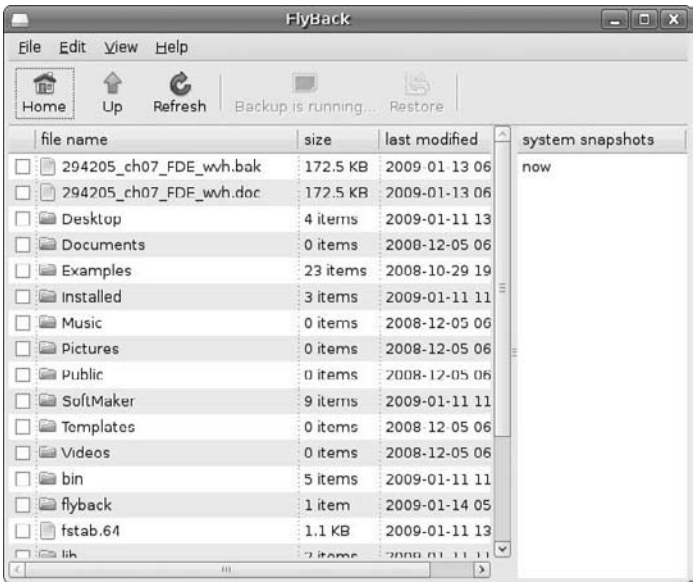
FIGURE B6-7

Back In Time's user interface and snapshot browser



FIGURE B6-8

Flyback's user interface and snapshot browser



- **TimeVault:** Under development by the folks at Canonical, TimeVault is GNOME-focused, integrated with Nautilus on Ubuntu systems, provides a great snapshot browser and restore interface, and is the solution that I'd recommend in Ubuntu-only environments. (I'm sure that Konqueror/Doplin integration is coming—many people have asked for it, and some have even volunteered to help.) Figure B6-9 shows TimeVault's snapshot browser, which makes it easy to locate and restore specific files or snapshots. Figure B6-10 shows Timevault's integration into Naultilus, which is available through its own tab in the Properties dialog for a file or directory. TimeVault is available as a Ubuntu/Kubuntu DEB package from its project page at `https://launchpad.net`. See `https://wiki.ubuntu.com/TimeVault` for more information about TimeVault.

At the time that this book was updated, none of these solutions was available through the Ubuntu repositories, but I'm sure that's just a matter of time.

FIGURE B6-9

TimeVault's snapshot browser

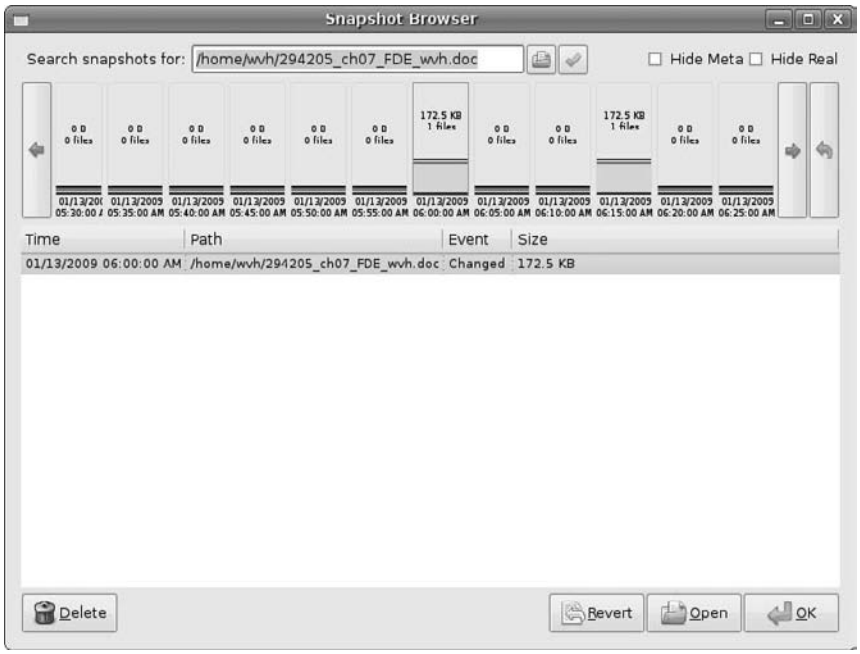
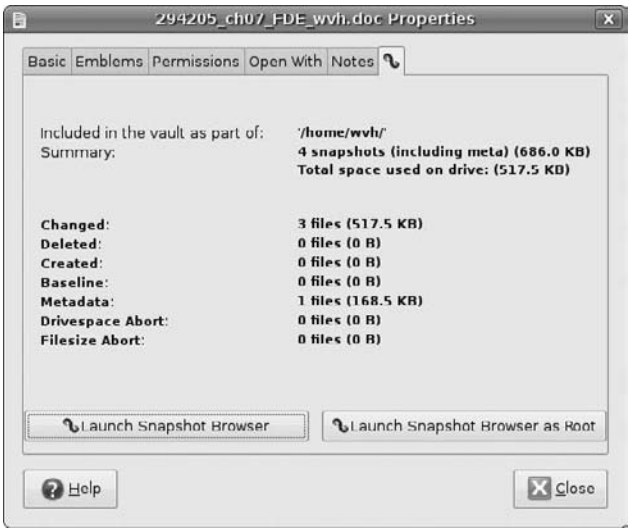


FIGURE B6-10

TimeVault's integration with Nautilus



Summary

This chapter discussed backups, which are the most time-consuming, boring, and valuable things that you need to do with your personal computer system. Ubuntu and Kubuntu are knee deep in excellent software that you can use to do various types of backups, ranging from backing up to files that you can store on removable media or simply move to other systems. It then discussed doing simple backups over the network and concluded by discussing BackupPC, an excellent open source backup application for doing regular, networked backups of multiple systems to a central backup server. The chapter concludes by highlighting some Linux backup packages that provide functionality similar to that of Apple's much-loved Time Machine backup/restore software.

Bonus Chapter 7 begins the server-oriented portion of the bonus material, discussing how to set up the Apache 2 Web Server on your Ubuntu system. Bonus Chapter 7 begins by providing a bit of background information about the Web and Web servers, and then explains the details of installing, configuring, and monitoring Apache 2, the latest and greatest Web server from the legendary Apache project.

Bonus Chapter 7

Using Virtual Machines and Emulators

Virtualization is the logical separation of the request for a service from the physical resources that actually provide that service, and is also one of today's hottest IT topics. The primary reasons for the popularity of virtualization today are the increasing speed and capabilities of today's off-the-shelf processors, falling hardware prices, and the ubiquity of networks that offer the capability to connect almost anything to anything else. These combine to provide higher performance and higher capability systems at a lower cost, providing plenty of CPU cycles and processor capabilities to handle any extra layer of abstraction that virtualization might introduce.

The loudest virtualization-related buzz surrounds different approaches to two topics, namely virtual storage and virtual machines:

- Virtual storage, as the name suggests, prevents users and system administrators from having to know the physical devices on which their data is stored because you access that storage through logical device, disk, or volume names. The physical hardware that supports these logical or virtual disks can be located on separate disk drives, storage devices, or even on physically distinct computer systems. Common virtual storage mechanisms such as logical volumes were introduced in Chapter 4 and discussed in more detail in Chapter 24.
- Virtual machines enable multiple complete operating systems to run simultaneously on the same physical computing hardware. Virtual machines must use the same instruction set as the system on which the virtual machines are running (which is typically referred to as the host system). In other words, you can only run x86-based operating systems on a machine with an x86-compatible processor. Virtual machines are often confused with emulators, which do approximately the same thing, but enable a single physical computer system to run operating systems and applications that require different processors and instruction sets by "emulating" those processors. Emulators are compiled for the instruction set of the machine on which they are running,

IN THIS CHAPTER

Introduction to Virtualization

Using VirtualBox

Using VMware Server

Using KVM

Using Xen

Using QEMU

but simulate an entire physical execution environment including the processor and hardware subsystems. Emulators running on x86-based systems can therefore run operating systems and applications that are compiled for either x86-based systems or for systems that use other types of processors, such as ARM, MIPS, PowerPC, and SPARC processors.

This chapter provides an introduction to different types of virtual machine technologies and discusses the different types of virtual machines and emulators that are available for Ubuntu and Kubuntu systems, usually as packages in the Ubuntu repositories.

NOTE **As mentioned previously, virtualization is a hot buzzword. To get as much software as possible on the bandwagon, you'll often hear remote desktop connection technologies such as VNC and RDP, discussed in Bonus Chapter 1, referred to as "desktop virtualization" technologies because they provide the ability to display a graphical desktop from one computer system on another computer system or smart display device.**

Overview of Virtualization and Emulation

As discussed in the introduction to this chapter, virtualization is a hot term in computing today. The next few sections explore some of the reasons to consider virtualization (using it as " résumé fodder" is not discussed), and the different types of virtual machines that are in use today (and which of these are available as packages for Ubuntu and Kubuntu systems), and make some recommendations on which virtual machine technology to adopt based on what you want to accomplish.

Why Virtualize?

The most obvious reason for wanting to run a virtual machine is so that you can run multiple operating systems at the same time on a single physical computer system. For example, suppose that you want to run Linux but occasionally need to run Microsoft Word. You can configure your system to dual-boot so that you can run either Linux or Windows, but that isn't very useful if you depend on your Linux system to read mail, browse the Web, and so on. Because you can run only one operating system at a time, dual-booting can be time-consuming and irritating. You could also run Windows using WINE on your Linux system—WINE is a Linux-based Windows API that supports Windows applications. However, you might encounter problems with WINE, and then what can you do?

Assuming that you have sufficient disk space, installing Windows as a virtual machine is probably your most convenient, robust, and long-term solution. Why waste hardware on Windows? Being able to run Windows while running Linux will enable you to use the Linux applications that you prefer while still being able to work on Microsoft Word documents that may have Windows-specific features (such as extensive Visual Basic macros, and so on).

Virtualization is more popular now than ever, for many reasons. Some of the most important reasons are the following:

- The power and performance of commodity x86 hardware continues to increase. Processors are faster than ever and support more memory than ever, and the latest multi-core processors literally enable single systems to perform multiple tasks simultaneously. Most modern systems can easily run a primary operating system and at least one virtual machine without taking a significant performance hit.

- The integration of direct support for hardware-level virtualization in the latest generations of Intel and AMD processors, motherboards, and related firmware makes virtualization on commodity hardware both easier and more powerful than ever before.
- A wide variety of virtualization products for both desktop and server systems running on commodity x86 hardware are available for free, and are easily installed on your Ubuntu or Kubuntu systems).

In general, the following are some common cases in which virtual machines are used:

- Dependency on a specific piece of software (described at the beginning of this section), where you need to use a software product that only runs (or runs best) on another operating system.
- Running legacy software, where you depend on a software product that only runs on a specific version of a specific operating system.
- Software system test and quality assurance environments, where you need to be able to test a specific software product on many different operating systems or versions of an operating system.
- Low-level development environments, where developers may want or need to work with specific versions of tools, an operating system kernel, and a specific operating system distribution.

Virtualization is not a new concept, and has been in use for years. However, for the reasons discussed in this section, virtualization is both more available and easier to run than ever before. The next section highlights different approaches to virtualization, and details which virtualization packages use them, and which of these are available for your Ubuntu or Kubuntu system.

Types of Virtual Machines

Virtual machines can work in many different ways. The differences between these approaches to virtualization can be subtle, but are always significant in terms of the capabilities that they provide and the hardware and software requirements for the underlying system. The most common approaches to virtual machines today are the following:

- **Guest OS:** Each virtual server runs as a separate operating system instance within a virtualization application that itself runs on an instance of a specific operating system. Parallels Workstation, VirtualBox, VMware Workstation, and VMware GSX Server are the most common examples of this approach to virtualization. The Ubuntu repositories include packages for the open source version of VirtualBox. A 15-day trial version of Parallels Workstation is available in the Ubuntu repositories. A free version of VMware Server can be downloaded and installed from the VMware Web site.
- **System-level or operating system virtualization:** Multiple, logically distinct system environments that are running on a single instance of an operating system kernel. System-level virtualization is based on the change root (chroot) concept that is available on all modern UNIX-like systems, and which gives each virtual machine a portion of the parent system's root filesystem as the root filesystem for the virtual machine, as well as its own set of processes that execute relative to that root filesystem. Linux virtualization solutions such as FreeVPS, Linux VServer, OpenVZ, and Virtuozzo are all examples of system-level virtualization. The kernel and user-space administrative utilities for OpenVZ are provided in the Ubuntu repositories.
- **Parallel virtual machine:** Some number of physical or virtual systems are organized into a single virtual machine using clustering software such as PVM (`www.csm.ornl.gov/pvm/pvm_home.html`). The resulting cluster is capable of performing complex CPU and data-intensive

calculations in a cooperative fashion. The PVM packages are available from the Ubuntu repositories, but are not discussed in more detail in this chapter because clustering is an advanced, specialized technology that requires its own book.

- **Paravirtualization:** Device drivers in the virtual machines are modified so that they "know" that they are running in a virtual machine environment, and can optimize device access for that model. This can significantly reduce execution overhead and increase disk and network I/O. Paravirtualization is typically used in virtualization environments that also use a virtual machine monitor (known as a hypervisor), which runs on top of your machine's hardware, manages protected or privileged CPU operations that are requested by any virtual machine, and also manages the hardware access requests that are made by any virtual machine. The most popular examples of hypervisor-based virtualization in use today are Citrix Xen, Microsoft's Hyper-V, and VMware's ESX and ESXi Server products, all of which run an administrative operating system on top of the hypervisor, as do the virtual machines themselves. This administrative operating system can communicate with the hypervisor and is used to manage the virtual machine instances and the resources that are allocated to them. Virtualization technologies such as KVM, where the Linux kernel that is running on the host system is equivalent to the hypervisor, are also actively working on paravirtualized device drivers.
- **Kernel-level virtualization:** Does not technically require a hypervisor, but instead runs a separate version of the Linux kernel and an associated virtual machine as a user-space process on the physical host. This provides an easy way to run multiple virtual machines on a single host, but only enables you to run Linux virtual machines, and then only those that have been compiled with support for the appropriate kernel-level virtualization technique. Examples of this are User-Mode Linux (UML), which has been supported in the mainline Linux kernel for quite a while, and KVM, which was introduced in the 2.6.20 mainline Linux kernel. UML does not require any separate administrative software in order to execute or manage its virtual machines, which can be executed from the Linux command line. KVM uses device drivers in the host's kernel for communication between the main Linux kernel and the virtual machines, and uses a slightly modified QEMU process as the display and execution container for its virtual machines. In many ways, KVM's kernel-level virtualization is a specialized version of Full virtualization, where the Linux kernel serves as the hypervisor, but I think that UML and KVM are unique enough to merit their own class of server virtualization. The Ubuntu and Kubuntu kernels include KVM modules, and the user-space component of KVM is available in the Ubuntu repositories.
- **Hardware virtualization:** Very similar to both paravirtualization and full virtualization, hardware virtualization uses a hypervisor, but is available only on systems that provide hardware support for virtualization. Hypervisor-based systems such as Xen, VMware ESX, and ESXi Server, and kernel-level virtualization technologies such as KVM can take advantage of the hardware support for virtualization that is provided on the latest generation of Intel (Intel VT, AKA Vanderpool) and AMD (AMD-V, AKA Pacifica) processors. Virtual machines in a hardware virtualization environment can run unmodified operating systems because the hypervisor can use the hardware's support for virtualization to handle privileged and protected operations and hardware access requests, and to communicate with and manage the virtual machines. Both Xen and KVM are available in the Ubuntu repositories.

Each of these approaches to virtualization has its own advantages and disadvantages. As mentioned previously, depending on your hardware, a single virtual machine technology can actually support more than one of these approaches to virtual machines. On Intel VT and AMD-V processors, Xen supports both paravirtualization and hardware virtualization. On these same processors, KVM supports both kernel-level and hardware virtualization.

What Kinds of Virtualization Can I Use?

With the exception of hardware virtualization, all of the virtualization techniques described in the previous section can be used on any x86-compatible system with a Pentium or better processor. I'd recommend using at least a Pentium IV or equivalent, and would also recommend using a 64-bit processor because those are modern, powerful, and often provide hardware support for vitualization.

If your system is using a processor that has hardware support for virtualization, such as the Intel VT and AMD SVM processors, you can run unmodified operating systems as virtual machines under virtualization technologies such as KVM and Xen. Hardware support for virtualization enables you to install and run operating systems such as Microsoft Windows by installing it in a virtual machine just as you would on a new physical system. (You can also, of course, use any vanilla Linux installation disk and install any standard Linux distribution as a virtual machine, without modification.) At the time of this writing, processors that provide hardware support for virtualization are as follows:

- **Intel VT (Virtualization Technology, aka Vanderpool):** Selected Pentium 4 and Pentium D, Xeon 5000 and later, Xeon LV, Core Duo, Core 2 Duo, Core 2 Quad, and Quad-Core processors.
- **AMD-V/SVM (Virtualization/Secure Virtual Machine, aka Pacifica):** AMD Athlon 64 (family "F" and "G") socket AM2, Athlon 64 X2 (family "F" and "G"), Turion 64 X2, newer Opteron, all Phenom, and all newer processors.

To determine whether the processor in a given system supports virtualization, you can check the flags entry in the file `/proc/cpuinfo`, which you can do with a command like the following:

```
cat /proc/cpuinfo | grep flags
```

Depending on whether you are using an AMD or Intel processor, the output of this command is very different:

```
AMD flags :   fpu tsc msr pae mce cx8 apic mtrr mca cmov  \
               pat pse36 clflush mmx fxsr sse sse2 syscall \
               nx mmxext fxsr_opt lm 3dnowext 3dnow up pni \
               cx16 lahf_lm svm cr8_legacy

Intel flags : fpu vme de pse tsc msr pae mce cx8 apic sep  \
               mtrr pge mca cmov pat pse36 clflush dts acpi \
               mmx fxsr sse sse2 ss ht tm pbe syscall nx lm \
               constant_tsc arch_perfmon pebs bts rep_good  \
               pni monitor ds_cpl vmx est tm2 ssse3 cx16    \
               xtpr lahf_lm
```

If you are using an Intel CPU, you would look for the field `vmx` in the flags field. If you are using an AMD CPU, you would look for the field `svm` in the flags field. The following command is an easy way to do this in one step:

```
# egrep '(vmx|svm)' /proc/cpuinfo
```

If you see any output from this command, your system supports hardware virtualization. Congrats!

Selecting a Virtualization Technology

The key to selecting a virtualization technology boils down to how much work you want to do and what you want to use the virtual machines for:

- Guest OS solutions such as VirtualBox and the free VMware Server are the easiest to install and understand, and they require limited changes to your system aside from installing them in the first place. The kernel drivers for VirtualBox are built into the standard Ubuntu and Kubuntu kernels, and VMware Server builds and installs its drivers as part of the installation process. Of these two, VirtualBox is the easier to use and manage because its kernel drivers will always be provided as part of any kernel updates, and you don't have to have a C compiler and related development tools installed on your system. However, VMware is more of an industry standard than VirtualBox, and any experience that you get using VMware Server may be useful both at home and at work.
- Hypervisor solutions such as Xen can provide higher performance and more flexible administration than other approaches virtualization, but require additional administrative overhead. Your system's boot configuration is different when booting a hypervisor and loading the domain Xen kernel. You must use specialized tools (xm, virt_manager, and so on) to configure and manage your virtual machines.
- Kernel-level virtual machines such as KVM are simpler to administer than hypervisor-based virtual machines, but do not provide the same level of administrative granularity of memory use, execution priority, and so on.

Recommendations? If you simply want to run the occasional Windows application, use VirtualBox. If you want to experiment with more sophisticated virtualization technologies from a desktop system, try KVM. If you want a full-blown, enterprise-ready virtualization environment and are willing to do a bit of learning, go for Xen. If you need to run virtual Windows machines and have an Intel VT or AMD-V processor, you can use either KVM or Xen to run virtual machines that take advantage of the hardware support for virtualization that is provided by these processors.

Popular Emulation Software for Linux

The best-known emulation package for most systems is not a desktop computer emulator at all—it's the Multi-Arcade Machine Emulator (MAME, at `www.mame.net`), which enables you to play a variety of ancient arcade video games directly from images of their ROMs. I love this package—arcade video games were my original motivation for getting interested in computers. Although system administration turned out to be more fun for me, playing some ancient video game via MAME is incredibly fun, and the MAME technology is simply amazing. There are some legal issues with playing many of the traditional game ROMs because they are still technically owned by the original manufacturers—you should investigate these before downloading any game ROMS from the hundreds of sites that offer them on the Web. However, the copyright to many of them has been lost, and I do actually have a Centipede arcade game in my living room, so I don't feel too bad about things. A variety of MAME packages are available in the Ubuntu archives—I typically just install the basic xmame-x package, which contains the binaries for the vanilla X Window system version of MAME. If you want to simplify the process of figuring out where to put your ROM images from the MAME documentation, front ends to this package such as gxmame (for GNOME systems) and kxmame (for KDE systems) can be very handy, although gxmame isn't available in the Ubuntu repositories at the moment. (You can get gxmame from its SourceForge home at `http://gxmame.sourceforge.net`.)

If you actually want to do real work with an emulator, the most useful emulation technology in Linux circles is QEMU, which can emulate many different hardware platforms and associated hardware. Not only is QEMU great as a standalone emulator, but open source virtualization packages such as Xen, VirtualBox, and KVM make heavy use of some of the code from QEMU to avoid re-inventing the wheel regarding console emulation, graphics emulation, and so on.

A huge number of emulators are available today, especially on Linux, where many people want to revisit the software glories of their youth (or at least the systems that they ran on). Rather than trying to coax a 20-year personal computer from a defunct vendor back to life, writing an emulator can be entertaining, is useful to others who are also interested in that hardware, and provides a fun technical challenge. A quick search for "Linux Emulators" using Clusty (`www.clusty.com`) or Google (`www.google.com`) returns hundreds of results. Two especially nice summary pages are Zophar's page of UNIX/Linux Emulators at `www.zophar.net/unix/unix.phtml`, and the Using Linux summary at `www.usinglinux.org/emulators/`.

Managing Virtual and Emulated Machines

As discussed later in this chapter, virtualization packages such as VirtualBox and VMware Server provide a graphical virtualization console that makes it easy to configure, manage, and switch between the virtual machines that you have configured and may be running using these technologies. However, these management consoles only work for virtual machines of the associated type. Many software packages are under development to provide a single management, creation, and execution framework for virtual machines that have been created with multiple virtualization technologies. The best known of these is the Virtual Machine Manager, and is available in the Ubuntu repositories as the virt-manager package. For more information about the Virtual Machine Manager, see its home page at `http://virt-manager.et.redhat.com`.

This chapter focuses on the manual configuration, creation, and management of different types of virtual machines rather than using a graphical management package. The primary reasons for this are the following:

- The Virtual Machine Manager only supports KVM, QEMU, and Xen virtual machines, but other virtualization technologies such as VirtualBox or VMware may be better suited to your needs.
- If you are planning to use a virtualization technology such as KVM or Xen in a server environment, a graphical user interface may not be available, in which case you cannot use the Virtual Machine Manager.
- Like most graphical tools, the Virtual Machine Manager is designed to simplify configuring, executing, managing, and monitoring its subject domain, in this case virtual machines. If you're just getting started with virtualization and plan to use it in a SOHO or enterprise environment, it's important to understand the system-level configuration changes and underlying procedures associated with virtualization, especially networking configuration.

The Virtual Machine Manager is a great package that can be extremely useful. However, in my opinion, understanding what's going on under the covers and focusing on the creation and management tools associated with specific virtualization technologies provides a better foundation for getting started with virtualization. My apologies if that disappoints you.

Some other virtualization management tools that are available in the Ubuntu repositories are the xenman package, for managing Xen virtual machines, and packages such as qemuctl, and qemu-launcher, for managing QEMU virtual machines.

Getting More Information About Virtualization

The Ubuntu Web site provides some general information about virtualization on Ubuntu, Canonical's partners in the commercial virtualization space, and so on. The sections of this chapter that discuss specific virtualization solutions will provide pointers to additional information about that particular solution. You can get additional, general information about virtualization on Ubuntu and Kubuntu systems at the following locations:

- **Ubuntu Virtualization Overview (`www.ubuntu.com/products/whatisubuntu/serveredition/technologies/virtualization`):** This is Canonical's corporate description of basic virtualization options for Ubuntu. Unfortunately, this doesn't mention things like VirtualBox and Xen! Luckily, the community documentation page that it points to does (the Ubuntu Community Virtualization Overview page in this list.)
- **Ubuntu Community Virtualization Overview (`https://help.ubuntu.com/community/Virtualisation`):** Ubuntu community documentation "home page" for information about different virtualization solutions on Ubuntu.

If you are just generally interested in virtualization and would like to obtain additional information, there are a variety of sites that provide general-purpose information about virtualization and emulation on commodity hardware. The following are some of my favorites:

- **IBM Developer Works:** (`www-128.ibm.com/developerworks/linux/library/l-linuxvirt`) This is a great article if you are interested in general information about virtualization on Linux systems. IBM's Developer Works always provides a huge amount of detailed, technically accurate information. This article provides a good overview and explanation of virtualization on Linux, along with links to many other sites.
- **Linux Virtualization Wiki:** (`http://virt.kernelnewbies.org`) A nice site that provides a good overview of virtualization on Linux and more detailed discussions of various Linux virtualization technologies. It also provides a good introduction to various tools that are available for managing Xen virtual machines, and a good overview of the Xen (and other) virtualization support that is available in various x86 Linux distributions (and in other operating systems, such as OpenSolaris).
- **SearchServerVirtualization.com:** (`http://searchservervirtualization.techtarget.com/`) A great TechTarget site that provides a tremendous amount of information about server virtualization, related storage virtualization topics, and so on. One of the advantages of this site is that many vendors make white papers and commercial announcements available there, which make it quite useful for keeping track of the pulse of commercial virtualization offerings.
- **TheFreeCountry.com:** (`www.thefreecountry.com/emulators/index.shtml`) A nice summary of free emulation and virtualization software for Intel x86–based (and compatible) machines. This is a great resource, but remember that free does not necessarily mean open source.
- **x86Virtualization Blog:** (`http://x86virtualization.com`) An interesting blog with some good technical insights. This site provides some good general information about topics such as using networks with virtual machines, desktop and enterprise virtualization, and so on. Unfortunately, Xen often seems to get short shrift here compared to VMware, Parallels, and SWsoft.
- **Wikipedia:** (`www.wikipedia.com`) Big surprise here. Wikipedia has extensive articles about Xen (`http://en.wikipedia.org/wiki/Xen`), virtualization in general (`http://en.wikipedia.org/wiki/Virtualization`), emulation (`http://en.wikipedia.org/wiki/Emulation`), and most of the other terms and players in today's virtualization market. The Wikipedia is a great source of information, as always—although you should always corroborate anything that you read there.

Installing and Using VirtualBox

VirtualBox (`www.virtualbox.org`) is an open-source Guest OS virtualization package that runs on and host's 32-bit and 64-bit operating systems such as Linux, *BSD, Solaris, and Microsoft Windows systems. VirtualBox itself runs on Linux, Mac OS X (Intel only), Solaris, and Microsoft Windows systems. VirtualBox is an open source application that was originally developed by InnoTek Systemberatung GmbH (`www.innotek.de/`), a software company based in Stuttgart, Germany, which was purchased by Sun Microsystems (`www.sun.com`). VirtualBox leverages the open source QEMU emulator for some of its basic console and input device support, as do other open source Linux virtualization packages.

VirtualBox comes in both a commercial and an Open Source Edition (OSE) releases. Precompiled commercial binaries are free for personal use or for evaluation purposes according to the terms of the VirtualBox Personal Use and Evaluation License (PUEL), while OSE sources and binaries are freely available under the terms of the GPL. Packages for the open source edition of VirtualBox are available in the Ubuntu repositories, along with the loadable kernel module that it uses (vboxdrv). This section discusses VirtualBox OSE, although the personal commercial version has some nice features that the OSE does not, such as the USB device and filesystem support, and the ability to seamlessly run applications from Windows virtual machines on your Linux desktop. (You can download the version of the commercial edition that is free for personal use from Sun at `www.sun.com/download/index.jsp`.)

For more information about VirtualBox, see:

- VirtualBox Home Page (`www.virtualbox.org/`): The home page for VirtualBox, which provides access to downloads of the latest (and beta) versions, the official VirtualBox documentation and related Wikis, and so on. A variety of forums where you can ask questions, exchange tips and tricks, and so on, are available at `http://forums.virtualbox.org/`.
- Sun's VirtualBox Page (`www.sun.com/software/products/virtualbox/index.jsp`): The page for VirtualBox at Sun Microsystems, the corporate overlords of commercial VirtualBox support.
- Community Ubuntu Documentation for VirtualBox (`https://help.ubuntu.com/community/VirtualBox`): Information about using VirtualBox on Ubuntu systems.

CAUTION **At the time of this writing, the OSE version of VirtualBox that was provided in the Ubuntu repositories is based on VirtualBox 1.5.6, while the PUEL version is 1.6.4. These use different (and incompatible) configuration file formats. If you install the OSE version of VirtualBox, create virtual machines, and then subsequently install the PUEL version (or any 1.6 version, or better), the PUEL version of VirtualBox will upgrade your configuration files and offer to make backups of your existing ones. You should definitely do this! If you need or want to revert to the OSE version of VirtualBox, you can then `cd` to the directory for each of your virtual machines (located in subdirectories of the `~/.VirtualBox/Machines` directory) and copy files with names like `machine-name.xml.1.2-linux.bak` to your actual configuration files, which have names like `machine-name.xml`.**

VirtualBox Overview

VirtualBox provides both command-line and graphical tools for creating, administering, and managing virtual machines. Its graphical administrative environment is much like those provided by tools such as Parallels Workstation, VMware Workstation and Server, and even Microsoft's Virtual PC, enabling you to create and customize virtual machine configurations.

The emulation and display environment used by VirtualBox is very similar to that used by VMware and others, although the key to release the cursor from being captured is different (and configurable). Internally, VirtualBox uses a kernel module called vboxdrv to communicate with and manage virtual machines and associated resources through the VirtualBox device node `/dev/vboxdrv`. This node must be publicly writable or all users who want to run VirtualBox must be members of the vboxusers group. (The latter is the default mechanism.) VirtualBox includes a built-in DHCP server and NAT support, to simplify connectivity between multiple machines, whether physical or virtual, while also making it possible to segregate and manage virtual machine network traffic.

VirtualBox is stable and quite responsive, and includes a number of nice features. One of these is the ability to specify the amount of video memory through a slider, which makes it possible to run more graphics-intense applications inside VirtualBox. Another is its built-in support for VirtualBox Remote Desktop Protocol (RDP), which enables you to use a standard RDP client to connect to virtual machines running inside VirtualBox (analogous to VMware's support for remote connections and administration). VirtualBox also supports virtual channels between local devices and RDP-connected virtual machines, enabling you to connect USB devices locally but attach them to remote virtual machines. VirtualBox also provides a nice snapshot feature, which makes it easy to capture the state of a virtual machine at any point, for easy rollback. This is great if you want to experiment with installing certain applications for short-term use, but don't want to do so permanently—just take a snapshot, install and experiment with the applications that you're curious about, and then fall back to a previous snapshot.

The next few sections explain how to create a virtual machine in VirtualBox OSE, install an operating system, configure networking for your virtual machines on both your host system and in the virtual machines, and start and stop virtual machines.

Installing VirtualBox

As mentioned previously, packages for the Open Source Edition (OSE) version of VirtualBox are provided in the Ubuntu repositories and can be installed on any Ubuntu or Kubuntu desktop system. VirtualBox requires a graphical user interface, and thus isn't really suitable for installation on a standard Ubuntu server system unless you've also installed a graphical user interface there. VirtualBox is an excellent virtualization package for desktop computer users, but isn't really oriented toward heavy-duty virtual machine deployment in an enterprise environment—Xen or VMware are better choices there, in my opinion.

To install VirtualBox on an Ubuntu or Kubuntu desktop system, use your favorite package management software (Synaptic on Ubuntu, Adept on Kubuntu, or the command-line tools anywhere), as described in Chapter 23. You will want to install the `virtualbox-ose` and `virtualbox-ose-modules-generic` packages. The former contains the actual VirtualBox binaries and related files, and the latter is a package that provides the loadable kernel module that is used by VirtualBox, but always depends on the latest module that is available for the latest kernel image. The latter will keep your VirtualBox installation working across Ubuntu or Kubuntu kernel upgrades.

NOTE **Installing VirtualBox as a specific user should automatically add that user to the vboxusers group. Members of this group have write access to the `/dev/vboxdrv` device that is used by VirtualBox for communication between the kernel and virtual machines. If you want other users of your Ubuntu or Kubuntu system to be able to use VirtualBox, you must manually add them to this group, which you can do with the command `sudo adduser USERNAME vboxusers`.**

Once VirtualBox is installed, the real fun begins. Proceed to the next section for information on creating a virtual machine in VirtualBox.

TIP **If you are going to be creating Ubuntu virtual machines, you can either follow the instructions in the next two sections for creating a virtual machine and installing an operating system in it, or you can use the `ubuntu-vm-builder` script to create a virtual machine (of type vbox) and automatically install Ubuntu into it. You must install the `ubuntu-vm-builder` in order to use this script, but it can be quite handy if you're oriented toward (or interested in) building virtual machines from the command line. After installing the package, see the man page for the `ubuntu-vm-builder` command or its online help (`ubuntu-vm-builder --help`) for more information.**

Creating a Virtual Machine in VirtualBox

You can start VirtualBox on an Ubuntu system by selecting the Applications ➪ System Tools ➪ VirtualBox OSE menu command, on a Kubuntu system by selecting the K Menu ➪ System ➪ VirtualBox OSE - PC Virtualization Solution menu command, or on either system by executing the `virtualbox` command from any command-line environment such as an xterm, GNOME Terminal, or Konsole window. When you first start VirtualBox, you'll see a screen like the one shown in Figure B7-1.

FIGURE B7-1

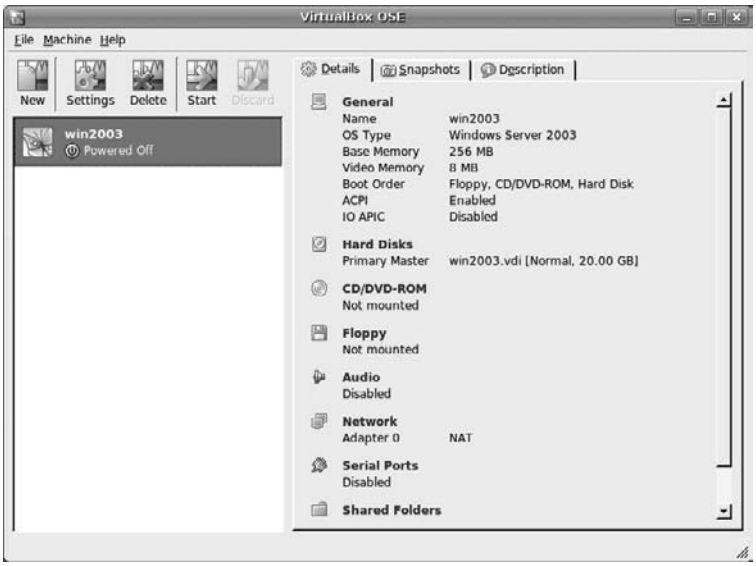
Starting VirtualBox for the first time



As the panel at the right of this screen states, this screen is largely empty because you haven't created any virtual machines yet. Once you create a virtual machine, the screen will look like the one shown in Figure B7-2, where the virtual machines that are available are listed in the left pane, and information about the currently selected virtual machine displays in the set of tabs at the right.

FIGURE B7-2

Starting VirtualBox after creating a virtual machine



To create a new virtual machine, click the New icon in the VirtualBox toolbar or select the Machine ➪ New menu command. A screen like the one shown in Figure B7-3 displays.

FIGURE B7-3

The Welcome screen for the New Virtual Machine Wizard

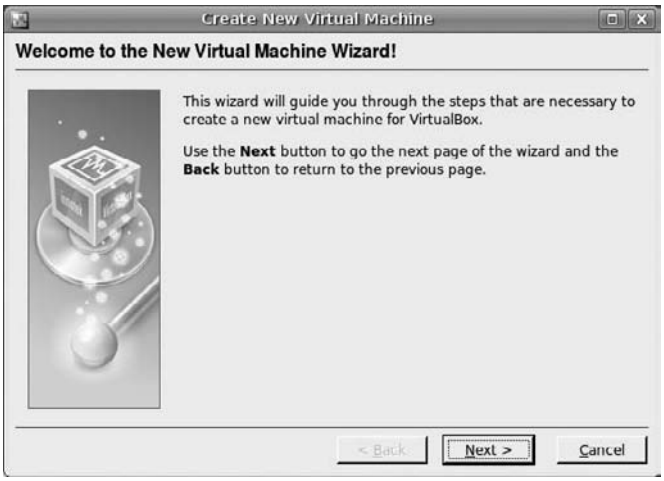
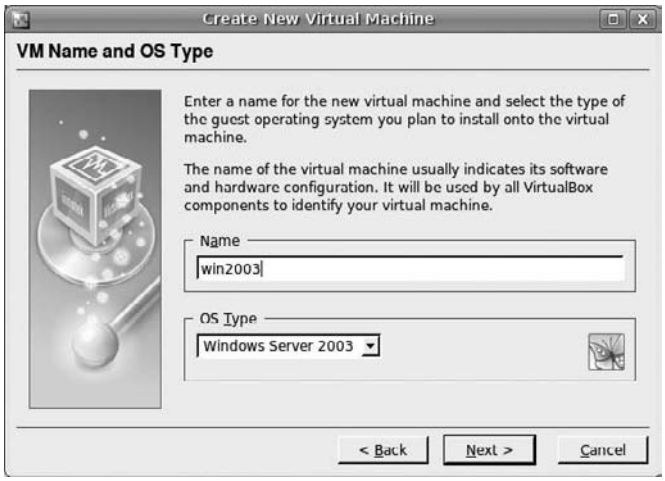


FIGURE B7-4

Providing name and type information for a new virtual machine

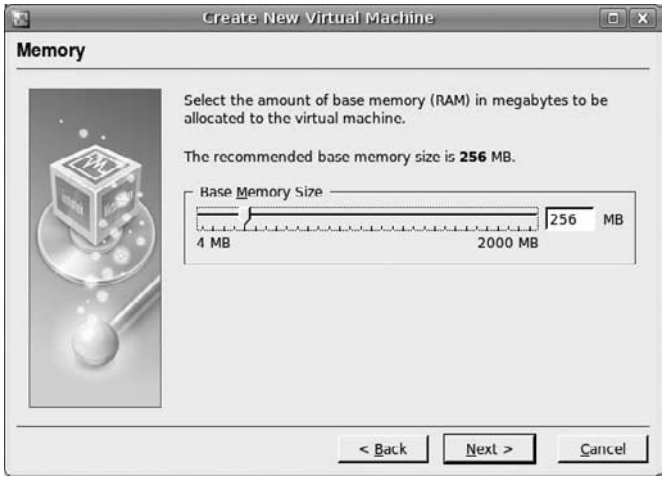


Click Next to proceed. A screen like the one shown in Figure B7-4 displays.

Enter the name that you want to associate with your new virtual machine in the Name field, and select the type of operating system that you are installing from the drop-down OS Type list. Click Next to proceed. A screen like the one shown in Figure B7-5 displays.

FIGURE B7-5

Specifying the amount of memory for a new virtual machine



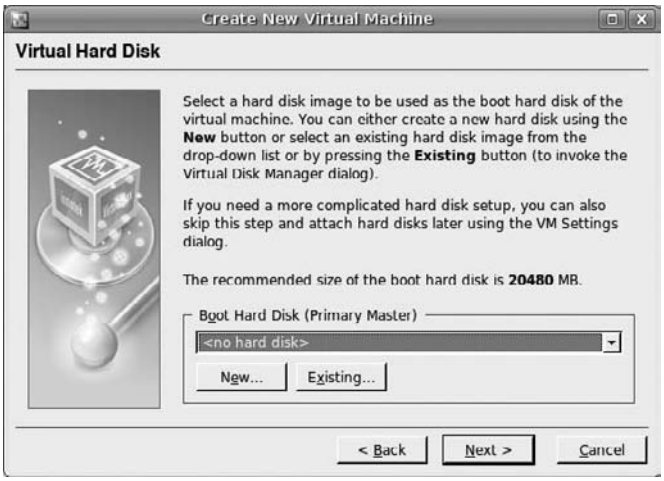
This screen enables you to specify the amount of memory that will be allocated to the virtual machine. This is the amount of physical memory that the virtual machine will believe that it has. You can change this later by modifying the settings for your virtual machine, so accepting the default suggestion (usually 256MB) is fine unless you know in advance that you will want more or less memory. For example, systems with graphical interfaces on which you will also be running server processes (such as a Web server, and so on) may need more memory. Systems without a graphical interface will usually require less memory, unless they are running heavily used server processes. As you might expect, any memory that you associate with a virtual machine will be allocated from the amount of memory available on the system where you are running VirtualBox, and will therefore be unavailable to the host system once the virtual machine is started. I typically accept the defaults and adjust them later if the performance of the virtual machine is unacceptable.

After optionally using the slider to change the amount of memory that will be allocated to your virtual machine, click Next to associate the specified amount of memory with your virtual machine and display the screen shown in Figure B7-6.

NOTE **The amount of memory that you specify on this screen will not actually be allocated until you start the virtual machine.**

FIGURE B7-6

Specifying the hard disk image for a new virtual machine



Because you are creating a new virtual machine, you will probably want to create a new virtual hard disk for its use. A virtual disk is a file on some storage device that the host system (the one on which you're running VirtualBox) can access. By default, new virtual disks are created in the `~/.VirtualBox/VDI` directory, located in your home directory, and have the file extension `.vmi`.

Click New to begin the Create New Virtual Disk Wizard, which makes it easy to create a new virtual disk. The Virtual Disk Wizard's Welcome screen displays, as shown in Figure B7-7.

NOTE **When creating a virtual machine in VirtualBox, you can also select an existing VirtualBox or VMware hard disk for its use by clicking Existing and navigating to the disk image in the Virtual Disk Manager that displays. In most cases you won't already have a disk image, so this section focuses on creating new virtual disks.**

Click Next to proceed. The dialog shown in Figure B7-8 displays.

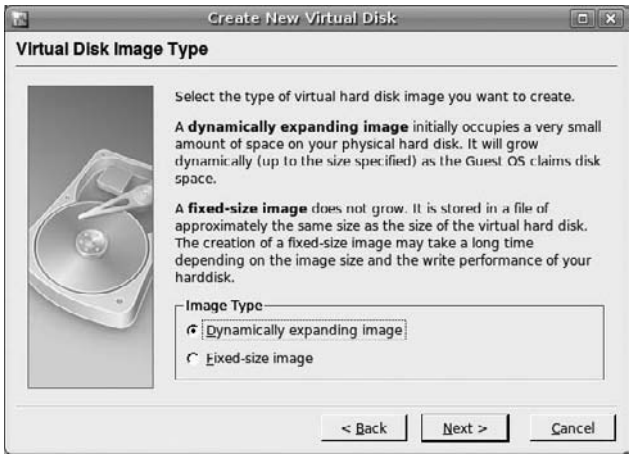
FIGURE B7-7

The Create New Virtual Disk Wizard



FIGURE B7-8

Specifying the type of virtual disk to create



This dialog provides two options for the type of virtual disk that you want to create, either a dynamically expanding image or a fixed-size image. The differences between these two are as follows:

- **Fixed-size image:** All of the disk space required on your host system for the virtual disk is immediately allocated when you create the new virtual disk, based on the disk size that you will specify in a subsequent dialog.
- **Dynamically expanding image:** Dynamically expanding images initially use the minimum amount of disk space on the host system that is required to hold the files that you've installed in your virtual machine, plus some extra space required to hold information about the filesystem(s) that the virtual disk image contains. For example, the 20GB dynamically expanding disk image for a Microsoft Windows 2003 Server system actually uses only 2GB after a complete, default system is initially installed. As you install additional applications and as the virtual machine creates and expands additional files for caching, logging, and so on, the dynamically expanding image will continue to grow, up to the maximum size that you will specify in a subsequent dialog.

These two different types of disk images have some performance implications for your virtual machine and the host system:

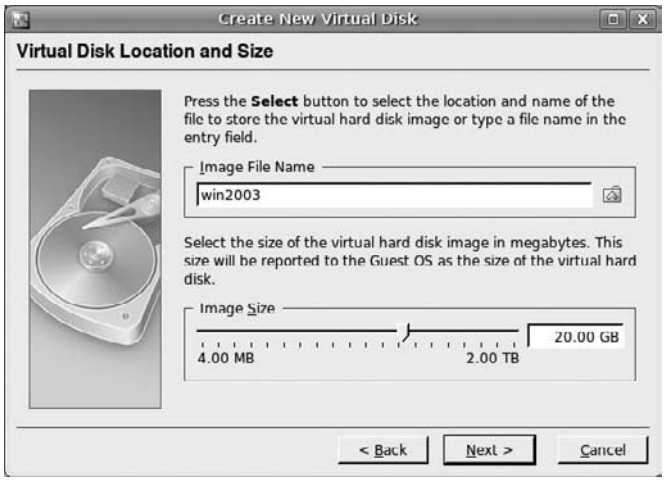
- **Fixed-size image:** Using a fixed-size disk image will result in slightly higher performance for your virtual machine and the host system because neither of them need to allocate additional space for the disk and do the housekeeping required to accurately track the disk blocks on the host system that are actually allocated to the virtual machine. If you know that you will be installing large applications or creating/downloading large files in a virtual machine, you may as well pre-allocate all of the disk space because you're probably going to need it eventually. On the down side, your host system will have less storage remaining after allocating the entire disk, and your backups and any virtual machine snapshots that you create will be larger because they must back up or snapshot the entire disk image.
- **Dynamically expanding image:** Using a dynamically expanding disk image results in slightly lower performance for both your host system and the virtual machine because they will have to dynamically allocate new space to the virtual disk image as it is required, and will also have to perform the housekeeping associated with that allocation. Your host system must always have enough disk space available for the maximum size of the virtual disk image—running out of disk space when trying to expand a virtual disk image can seriously corrupt the virtual disk image, which could render the virtual disk unusable, meaning that you could lose all of the data that it contains. On the plus side, dynamically expanding images require substantially less storage on the host system, consume less space in your backups, and will require less space for any snapshots that you create of your virtual machines.

I recommend using a dynamically expanding image because it requires less space for the disk itself, snapshots of the virtual machine, and in the backups that you do of your host system. (You are doing backups, right?). The performance impact of using a dynamically expanding disk is relatively small, and I only use VirtualBox to enable me to run the occasional Windows application or for testing other Linux distributions, neither of which is a situation in which optimal performance is critical. As I've mentioned before, I typically use Xen for mission-critical virtual machines where performance is critical. Your mileage may vary.

Click the radio button beside the type of disk image that you want to create and click Next to proceed. The dialog shown in Figure B7-9 displays.

FIGURE B7-9

Specifying the name and size of a new virtual disk



The name of a virtual disk image is initially set to the name that you specified for your virtual machine, but you can change this by entering a new name if you want. I rarely change the name of the disk image because using the same names for both makes it easy to see which disk image goes with which virtual machine if you explore the `~/.VirtualBox` folder. The default size of the disk image for a virtual machine is typically set to 8GB, which is enough for experimentation but probably not if you plan to install a large number of applications and/or do a significant amount of real work in your virtual machine. For example, I tend to use 20GB disk images for Windows systems to reduce the chance that I'll run out of space some day (and because disk space is cheap nowadays). If you're using dynamically expanding images, you won't lose all of that disk space on your host system immediately, although you should really have that much available just in case the virtual machine actually needs it.

After making any changes that you want to the name or maximum size of your virtual disk image, click next to proceed. A Summary dialog like the one shown in Figure B7-10 displays.

Click Finish to proceed. The dialog shown in Figure B7-6 redisplays, this time displaying the name and location of your new virtual disk as the Boot Hard Disk for your system. Click Next to proceed. The Summary dialog shown in Figure B7-11 displays.

This dialog summarizes the name, type, memory allocation, and virtual disk image for the new virtual machine that you are creating. Assuming that you've specified the values that you actually want to use in the preceding dialogs, nothing on this dialog should be a big surprise. You can change these settings later anyway, so you may as well proceed. Click Finish to proceed. The main VirtualBox window redisplays, showing an entry for your new virtual machine in the left column.

At this point, you've created a virtual machine, but it doesn't actually contain an operating system or any applications. If you're using a dynamically expanding disk, its virtual disk image has been created, but is using less than 100KB of storage. To actually run the virtual machine, you must install your operating system of choice, as described in the next section.

FIGURE B7-10

Summary information about a new virtual disk

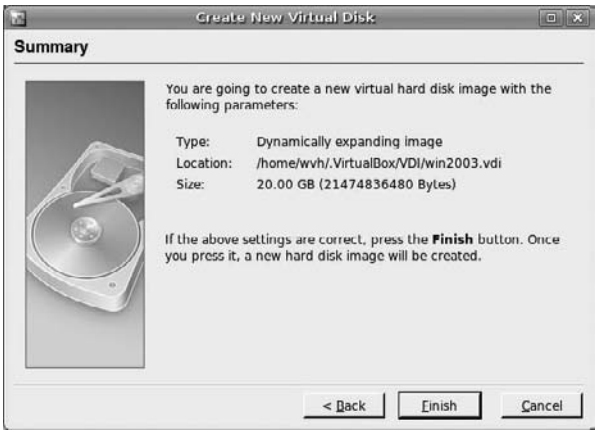


FIGURE B7-11

Summary information about a new virtual machine



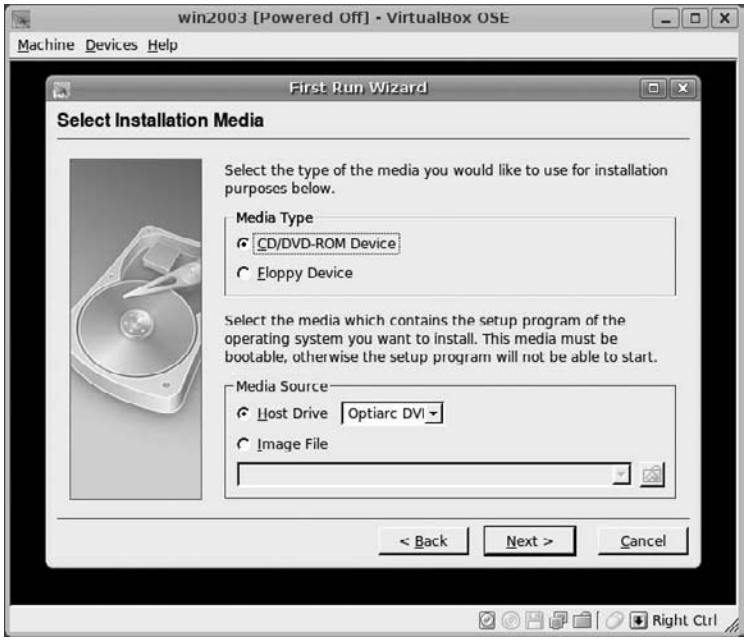
Installing an Operating System in a Virtual Machine

VirtualBox provides a helpful First Run Wizard to walk you through the steps of installing an operating system in a new virtual machine. To install your operating system of choice into a new virtual machine, you simply start the virtual machine by making sure that its entry is selected in the left-hand VirtualBox pane and then clicking the Start button above that pane. This displays a VirtualBox window that is the console for

your new virtual machine, and also displays the welcome screen for the First Run Wizard, which isn't very interesting. To get to a useful screen, click Next. This displays the dialog shown in Figure B7-12.

FIGURE B7-12

Specifying the installation media for a new virtual machine



This dialog enables you to boot from a CD/DVD-ROM drive or a floppy disk to start the installation process for the operating system for your virtual machine. After Selecting either of these as the media type, you can identify the source of that media, either a physical device or an image file in the bottom of the dialog. If you are booting from an image file rather than a physical device, you can click the browse button beside the Image File entry to browse for the location of your disk or floppy image. Once you've selected the Media Type and Media Source, click Next to proceed. The Summary dialog shown in Figure B7-13 displays.

Beyond simply summarizing the boot information for your new virtual machine, this dialog points out that you may need to use the Devices ➪ Unmount command to manually unmount your boot media if a reboot occurs during installation. This is a good thing to remember if you find that your virtual machine is booting from the installation media again after a reboot.

To proceed, click Finish. The First Run dialog closes, and your new virtual machine should start to boot from the specified installation media.

To interact with the virtual machine during the installation process (and beyond), you will need to click inside the VirtualBox window so that the virtual machine has the input focus for your mouse and keyboard. The first time you do that, you will see a dialog like the one shown in Figure B7-14.

FIGURE B7-13

The installation media summary for a new virtual machine

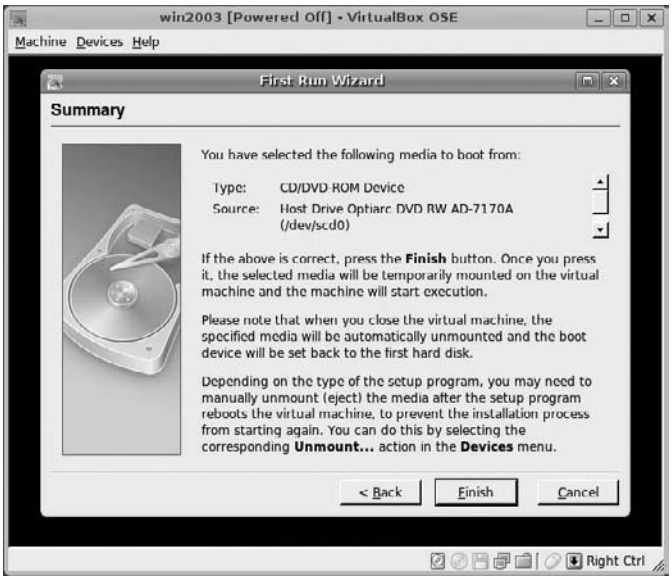
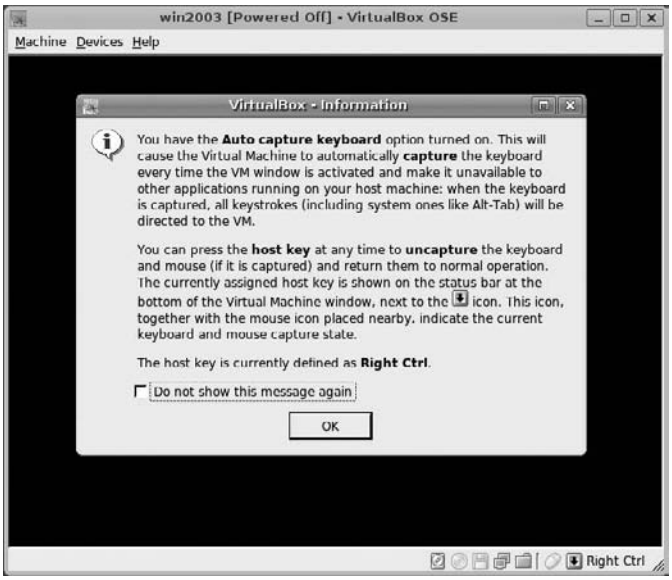


FIGURE B7-14

VirtualBox keyboard/mouse focus explanation dialog



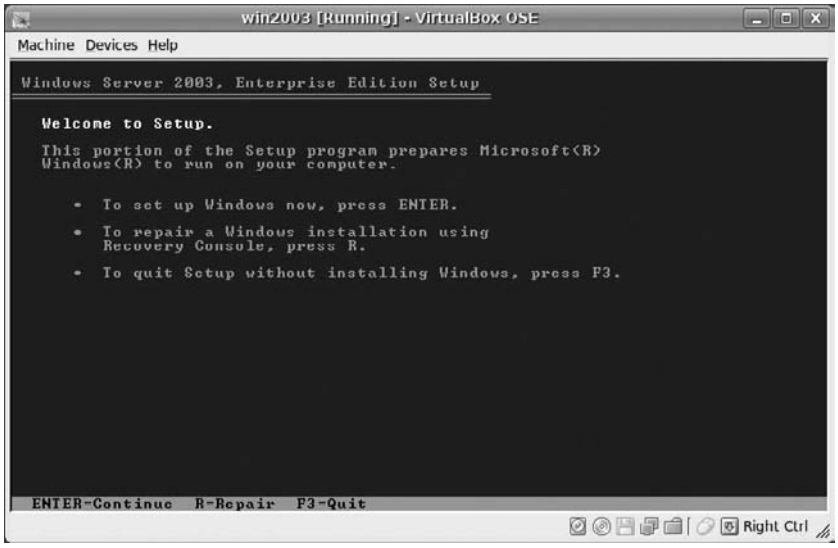
The Auto capture keyboard option makes it easy to interact with a virtual machine by sending all mouse and keyboard events directly to the virtual machine rather than the host system. This makes it possible to type keyboard commands such as the Ctrl+Alt+Delete required for Windows logins that might otherwise accidentally shut down your host system, and makes interacting with the virtual machine more precise because the cursor can't leave the boundary of the virtual machine window while you're in this mode. To release the mouse and cursor from your virtual machine, press the right Control (Ctrl) key. (You can change this key binding later if you want.) As a handy reminder, the key to release keyboard capture is also displayed in the lower-right corner of a VirtualBox virtual machine window.

After the first time you see this warning dialog, you'll probably want to disable it in the future, which you can do by checking the Do not show this message again option before clicking OK to close this dialog.

At this point, the installation of your virtual machine begins, using the standard installation screens for the operating system that you are installing, just as if you were installing the operating system on a physical machine, as shown in Figure B7-15.

FIGURE B7-15

Installing an operating system in a virtual machine



Once the installation process completes, your virtual machine will typically reboot. If necessary, remove the installation media as discussed earlier in this section and shown in Figure B7-13. You can now use your new virtual machine—congratulations!

TIP **If you are already using virtual machines under some VMware product and would like to use those same virtual machines with VirtualBox, you can convert them from VMware's virtual disk format (VMDK) to VirtualBox's virtual disk format (VMI) by following the instructions provided on the Ubuntu Wiki at** `https://wiki.ubuntu.com/UbuntuMagazine/HowTo/Switching_From_VMware_To_VirtualBox:_.vmdk_To_.vdi_Using_Qemu_+_VdiTool`.

Installing VirtualBox Guest Additions

The open source edition of VirtualBox provides enhanced drivers for better video resolution and better control of mouse interaction in Windows and Linux virtual machines once a virtual machine of either of those types has been installed. These drivers are known as the VirtualBox Guest Additions—if you're familiar with VMware, these are equivalent to VMware's Guest Tools.

NOTE **The Windows drivers in the Guest Additions package are pre-compiled. Installing the Linux drivers from this package requires that the GNU C Compiler (gcc) and the header files for the Linux kernel be installed in your virtual machine. See Bonus Chapter 3, "Software Development on Ubuntu," for information about installing these in an Ubuntu or Kubuntu virtual machine (which is the same as installing them on a physical Ubuntu or Kubuntu system). If you're running some other Linux distribution in your virtual machine, use that system's package management tools to install the `gcc` and `kernel headers` packages for that distribution.**

To install the Guest Additions into a virtual machine, select the Devices ➪ Install Guest Additions menu entry in a virtual machine window. You will see a dialog to the effect that the Guest Additions CD image couldn't be found on your machine, and would you like to download it over the Internet? Click Yes. If you see a confirmation dialog for the download ("are you sure"), you should also click Yes there, and you will see a download progress bar appear in the bottom-left corner of the virtual machine window as the image is downloaded.

Once the ISO image for the Guest Additions is downloaded, VirtualBox displays a dialog asking if you want to mount the ISO image as a CD drive. Click Yes.

If your virtual machine is running some version of Windows, the Autorun.inf script on the CD will automatically install the Guest Additions for you when the CD is mounted. If you are running Linux, you will have to start a terminal window (xterm, GNOME Terminal, Konsole, and so on), `cd` to the mounted image, and run the `VBoxLinuxAdditions.run` script via `sudo` to build and install the enhanced drivers.

Once you've installed the drivers, you can unmount the ISO image by using the standard mechanism for your operating system, or by selecting the Devices ➪ Unmount CD/DVD-ROM menu command from the virtual machine window.

Starting and Stopping Virtual Machines

Once you've created one or more virtual machines and installed an operating system in each, you can start any virtual machine by selecting its entry in the left pane of the VirtualBox window and clicking the Start icon.

After you're finished using a virtual machine, you can shut it down in one of two ways:

- By invoking the standard shutdown process for the virtual operating system that your virtual machine is running. When the shutdown process completes, the virtual machine window closes and only the standard VirtualBox dialog will be displayed (unless you are running other virtual machines in other virtual machine windows).
- By clicking the Close box in the upper-right corner of the virtual machine window or selecting the Machine ➪ Close menu command from the virtual machine window. Both of these display the dialog shown in Figure B7-16, which enables you to shut the virtual machine down in various ways.

FIGURE B7-16

Different ways of closing a virtual machine



The three options in this dialog do the following:

- **Save the machine state:** Saves the memory and disk state of the virtual machine to a file that will be used the next time that you start the virtual machine. Starting from a saved state is much faster than actually booting a virtual machine. However, you cannot edit the settings for a virtual machine if you have a saved state file because the state information for the virtual machine depends on the settings that were active when the machine's state was saved. In order to modify any settings for a virtual machine that has a saved state file, you must either start the virtual machine from the saved state file and then shut it down normally, or you must discard the saved state information by right-clicking on the virtual machine in the left pane of the main VirtualBox window and then selecting the Discard command from the context-sensitive menu. If you discard saved state information, you will lose any changes that were not written to your virtual machine's disk. If you have the time, it's definitely a better idea to start the virtual machine and shut it down normally rather than simply punting the saved state file.
- **Send the shutdown signal:** Begins a standard shutdown in the virtual machine's operating system. This option is provided if you accidentally clicked Close but have the time to do an orderly shutdown.
- **Power off the machine:** Immediately turns off the virtual machine without doing a clean shutdown. This is usually a bad idea because all operating systems need to ensure internal consistency before shutting down, but may be necessary in emergencies. However, if you've taken snapshots (described later in this section), the dialog shown in Figure B7-16 provides another option, shown in Figure B7-17, which enables you to fall back to the current snapshot, using it as a saved state file for this virtual machine. (Don't worry, I'll talk about snapshots in a few more paragraphs!)

FIGURE B7-17

Reverting to a snapshot during power-off



To proceed, select the way in which you want to shut down the virtual machine, and click OK.

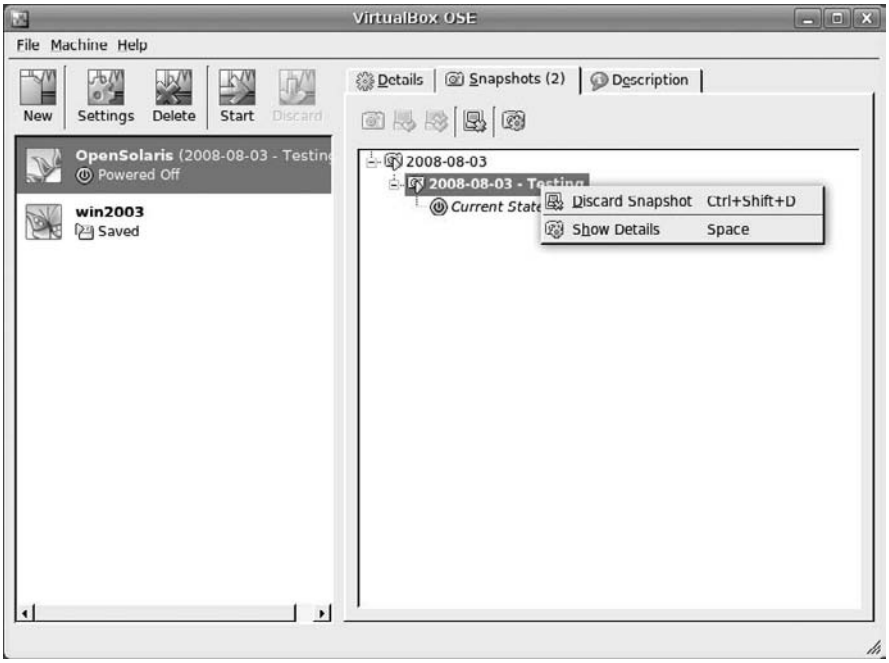
In general, I prefer saving the state of a virtual machine to the other virtual machine shutdown options. However, these files can use significant amounts of disk space, and can therefore require significant amounts of space in your backups (as well as increasing the time that it takes to do backups). You are doing backups, right?

A interesting capability that virtual machines introduce is the ability to take a snapshot of the state of a virtual machine and then undo any changes that you've made to your system by falling back to the previous state of the virtual machine. VirtualBox snapshots work slightly differently than you might think. When you create a snapshot using the Machine ➪ Take Snapshot menu command from a virtual machine window, you actually continue to run in the snapshot. Once you've taken a snapshot, you can fall back to any previous snapshot by selecting the Snapshots tab for that virtual machine in the main VirtualBox window, right-clicking on the snapshot that contains the changes that you want to undo, and selecting the Discard Snapshot command from the context-sensitive menu that displays, as shown in Figure B7-18.

Selecting this command deletes the current snapshot. The next time that you start your virtual machine, it will start with the previous snapshot or, if no snapshot is available, with the default state of your virtual machine.

FIGURE B7-18

Deleting a snapshot to return to a previous state



Networking and VirtualBox

VirtualBox includes an internal DHCP server and built-in support for NAT (Network Address Translation) that enables your virtual machines to automatically obtain IP addresses on a private, Internet subnet. Virtual machines that you create in VirtualBox automatically use NAT networking to reach external networks such as the Internet. Unfortunately, such communications are handled by VirtualBox rather than through standard networking infrastructure. This means that, by default, you can't use standard network commands from your host system to reach your virtual machines, and you also can't reach the host machine from your virtual machines. Bogus! Luckily, VirtualBox supports other networking mechanisms, specifically Host Interface networking, that eliminate these hassles by enabling your virtual machines to communicate with other hosts and networks by using a virtual network interface on the host system, rather than through VirtualBox itself. This enables communication between your virtual machines and any other system, including your host system, which also enables any other system that can reach your host system to reach your virtual machines. As a plus, it isn't hard at all to set this up.

TIP **You can also use port mapping to map ports on your host system to the ports associated with various services on your virtual machine, which will enable other systems to contact your virtual machine, but this doesn't help with access from the virtual machine to your host system. But, it's better than nothing. For more information about port mapping in VirtualBox, see the sidebar "Remapping Ports for VM Access" at the end of this section.**

VirtualBox's Host Interface support uses a technique known as bridging that is implemented in software on your Linux host. Bridging is a networking technique used to connect multiple network segments by forwarding packets from one to the other. In this context, bridging enables you to enable communication between multiple physical and virtual network interfaces by connecting them to a software bridge that you set up on your Ubuntu or Kubuntu system. The network devices that are attached to this bridge consist of your system's primary network interface and one virtual network interface for each of your virtual machines. The bridge therefore sees any network traffic that the physical interface that is attached to it would originally have seen, and can route traffic to any other device on the bridge. This includes both connection traffic and infrastructure-related traffic such as DHCP requests, which therefore enables your virtual machines to get IP configuration information from a DHCP server on your local network, giving them standard addresses, standard DNS configurations, and any other bells and whistles that you've set up in your network's DHCP server. Of course, you can also give your virtual machines static addresses on your local network and supply other network configuration information as you would on any other host with a static IP address. It's all up to you!

NOTE **Bridging is not supported by many of the drivers for wireless network interfaces. If you want to use virtual machines on a laptop that uses (or only has) a wireless interface and need to be able to communicate between your laptop and your virtual machine(s), it's easier to install and use VMware Server than it is VirtualBox. You can selectively remap selected ports to enable interaction between a VirtualBox virtual machine that uses NAT and its host machine, but this is a pain. (See the sidebar at the end of this section for specific information about how to do that.)**

Thanks to the cleverness of the Linux, Ubuntu, and VirtualBox folks, setting up a bridge and virtual network interfaces on your host system does not require a certificate from Cisco, the ability to recite the seven layers of the OSI networking model, or the sacrifice of any form of poultry. Let's just do it! I'll present these instructions as a series of steps to make it easier to follow along (and to highlight the fact that there are not that many steps):

1. Use your favorite package management application to install the `bridge-utils` and `uml-utilities` packages on your system. The `bridge-utils` package contains the commands that are used to create bridges, attach devices to bridges, show status information, and so on. The

uml-utilities package contains the utilities that are used to create and manage Linux virtual Point-to-Point (TUN) and Ethernet (TAP) network devices. (The "uml" stands for User-Mode Linux, which is a mechanism for running instances of the Linux operating system as a process on a Linux machine, which was mentioned earlier in the section "Overview of Virtualization and Emulation" of this chapter.) For example, you could install these packages using aptitude from a terminal window by using the following command:

```
sudo aptitude -y install bridge-utils uml-utilities
```

2. Modify your /etc/network/interfaces file to add an entry for a bridge and connect the network device that is associated with your system's physical network interface to it. You should also comment out any existing entry for that network device. The following is a sample entry for a bridge named br0 to which the network device eth0 is attached:

```
iface br0 inet dhcp
  bridge_ports eth0
  bridge_fd 0
  bridge_maxwait 0
  metric 0
```

This entry says that the br0 is an IP bridge (because of the bridge_ports entry) that should get its configuration via DHCP, that no delay should be imposed when forwarding packets across the bridge (the bridge_fd entry), that the bridge should immediately enter forwarding mode as soon as it comes up (the bridge_maxwait entry), and that this interface should be considered a high-priority interface that can immediately reach a default gateway (the metric entry). Remember to comment out any existing entries for your network device (eth0, in the example used in this section)!

TIP **Unfortunately, Ubuntu and Kubuntu systems don't provide a man page that describes the settings that you can use to define bridge settings in an** /etc/network/interfaces **file. For more information about the settings I've used (and the many others that I did not), see the man page for the** bridge-utils-interface **additions, which you can find on the Net at locations such as** http://web.iesrodeira.com/cgi-bin/man/man2html?bridge-utils-interfaces+5**.**

3. Add an entry for a virtual network interface for each of your virtual machines to the /etc/vbox/interfaces file. These virtual network devices will be automatically created by the VirtualBox startup script (/etc/init.d/virtualbox-ose for the Open Source Edition of VirtualBox) when it is run during your system's bootup process. These entries look like the following:

```
tap1 wvh br0
tap2 wvh br0
```

The first field is the name of a virtual network device that you want to create, the second is the name of the user with which that network device will be associated, and the third is the name of the bridge that you want to attach the device to. These sample entries will create the virtual network devices tap1 and tap2, owned by the user wvh, and will attach them to the bridge named br0. You should add enough entries to this file so that there is a one-to-one correspondence between the number of virtual machines on your system and the number of TAP device entries in this file. Each TAP entry must have a unique name.

4. Modify the rule file that is used when devices associated with VirtualBox are created during your system's boot process, adding a rule that specifies that the TUN/TAP virtual network device `/dev/net/tun` will be writable by anyone in the vboxusers group. You will want to add the following entry to the end of the `/etc/udev/rules.d/50-virtualbox-ose`.rules file. The following should be entered on a single line, without the `\` character:

```
KERNEL=="tun", NAME="net/%k", OWNER="root", \
     GROUP="vboxusers", MODE="0660"
```

5. If you're using an Ubuntu system, disable the GNOME network manager. I'm a big fan of this tool in standard networking situations, but it just doesn't understand bridging. Most operating systems that support TCP/IP use a routing table to identify how to reach other hosts on various networks, including an entry for the default route that should be used when connecting with systems that aren't on your local network. Unfortunately, the network-manager insists on also creating routes for the first physical Ethernet device that it finds on a system, including a default route. This can lead to routing confusion, as in the following example:

```
$ route -n
Kernel IP routing table
Destination Gateway    Genmask         Flags Metric Ref Use Iface
192.168.6.0 0.0.0.0    255.255.255.0 U    0      0      0  br0
192.168.6.0 0.0.0.0    255.255.255.0 U    0      0      0  eth0
0.0.0.0     192.168.6.1 0.0.0.0      UG   0      0      0  eth0
0.0.0.0     192.168.6.1 0.0.0.0      UG   0      0      0  br0
```

As you can see, this system has two sets of routes, one via the bridge (which is fine) and one via `eth0` (which is on the bridge). This is one of those cases where two of something is not better than one. The duplicated default route can cause packets sent to any address that is not on the local network to loop in the `eth0` device and never cross the bridge. To quickly eliminate this problem, you can simply disable the network-manager. (You can also simply remove the network-manager package, but who knows if you may want to use it eventually—it might be fixed to deal with bridging, after all.) There are many ways to disable the network-manager, the easiest of which is to immediately cause it and its dispatcher subprocess to exit when they are first started. To do this, execute the following two commands:

```
sudo echo "exit" > /etc/default/NetworkManager
sudo echo "exit" > /etc/default/NetworkManagerDispatcher
```

The next time you start your system (which we'll do later in this section), no more network manager! This approach to disabling it is also easy to undo, by deleting these two files and then restarting the `dbus` command script using the command `/etc/init.d/dbus restart`.

6. Reboot your system. When your system comes back up, you can use the `brctl show` command to verify that a bridge has been created and that the right physical and virtual network devices are attached to it. The output associated with the example that I've used in this section looks like the following:

```
$ brctl show
bridge name bridge id           STP enabled interfaces
```

```
br0         8000.0019d1239f01   no          eth0
                                            tap1
                                            tap2
```

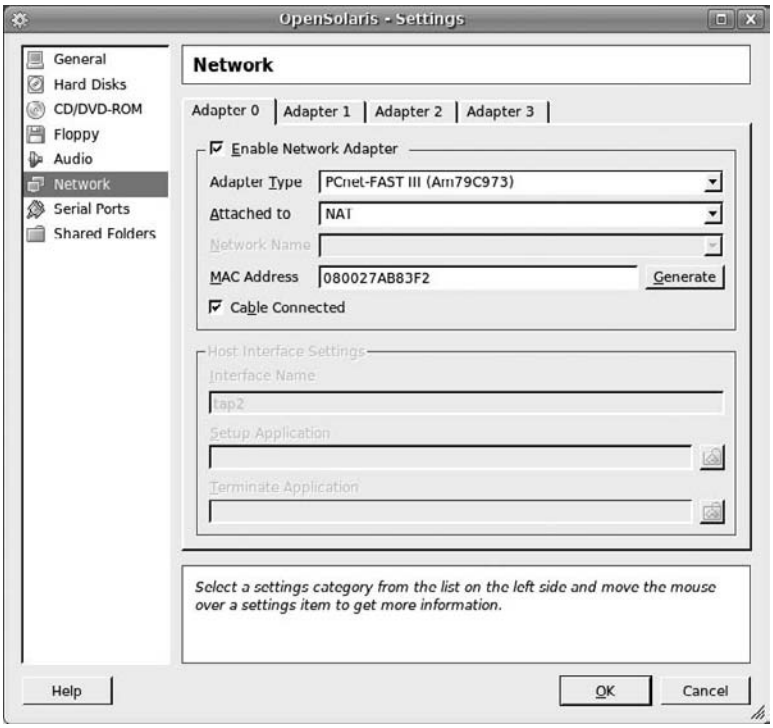
You should also verify that routing is set up correctly by executing the `route -n` command—you should see only a few entries, all related to the bridge.

7. Start VirtualBox and update the network configuration information for each of your virtual machines. To do this:

a. Select a virtual machine in the left pane, and click the Settings icon in the VirtualBox toolbar. The Settings dialog displays, showing the current configuration of the specified virtual machine. Click the Network entry in the left pane to display a dialog like the one shown in Figure B7-19.

FIGURE B7-19

Default network settings for a virtual machine

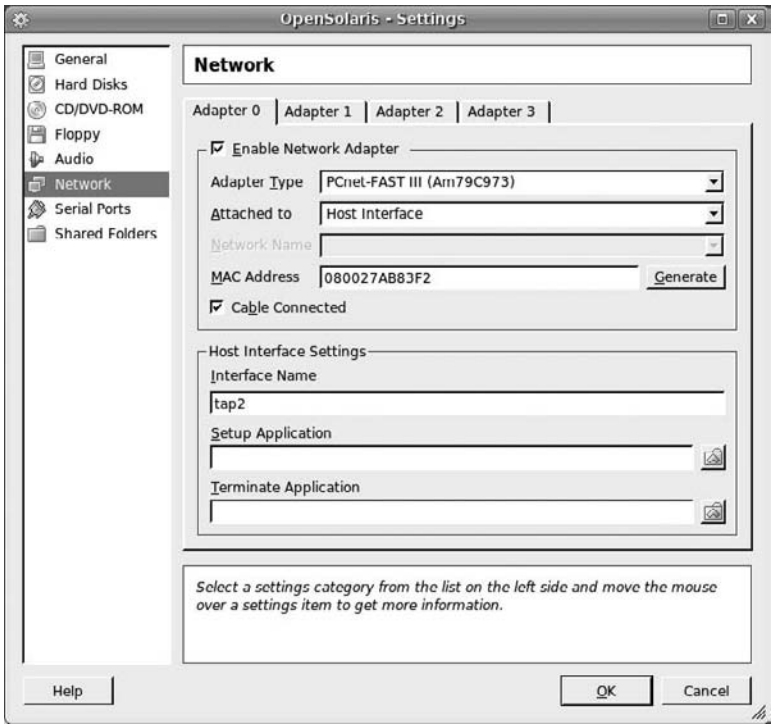


NOTE **The Settings icon is available only for systems that have been shut down cleanly and which do not have snapshots or saved state information. You must delete your snapshots and saved state information before you can modify any of the settings for a virtual machine.**

b. Click the down arrow beside the Attached to entry and select Host Interface from the drop-down list. This enables the fields at the bottom of the Network dialog. In the Interface Name field, enter the name of one of the TAP devices for which you created an entry in the `/etc/vbox/interfaces` file, as shown in Figure B7-20.

FIGURE B7-20

Host interface setup for a virtual machine



Each virtual machine must use a unique TAP device.

c. Click OK to save the updated network configuration information.

After updating all of your virtual machines, any virtual machine that you start will now be a true citizen of your local network. Not too shabby!

Remapping Ports for VirtualBox VM Access

Port forwarding is a technique where one system forwards requests on certain network ports to another system. Nowadays, port forwarding is probably most commonly used on home gateways to forward requests from an external gateway to a specific system on an internal network. VirtualBox provides built-in support for port forwarding to enable external systems (or your host system) to contact services on your virtual machine directly, without having to make the virtual machine a full-fledged participant in your local network.

You use the `VBoxManage` command to set up port forwarding, which must be done when the virtual machine is not actually running. Doing this requires three commands for each service that you want to be able to contact on the virtual machine, as in the following example, which enables external access to an SSH server that is running on a virtual machine named Ubuntu-Hardy:

```
VBoxManage setextradata "Ubuntu-Hardy" \
    "VBoxInternal/Devices/pcnet/0/LUN#0/Config/ssh/Protocol" TCP
VBoxManage setextradata "Ubuntu-Hardy" \
    "VBoxInternal/Devices/pcnet/0/LUN#0/Config/ssh/GuestPort" 22
VBoxManage setextradata "Ubuntu-Hardy" \
     "VBoxInternal/Devices/pcnet/0/LUN#0/Config/ssh/HostPort" 2222
```

(These commands would normally be typed on a single line—the backslash in each is present only for formatting purposes in this book.)

The first of these commands identifies the networking protocol that you are forwarding. The second identifies the port on the virtual machine to which requests on a specific port on your host should be forwarded. The third identifies the port on your host system that external systems should use in order to contact the service associated with this rule. The name of this rule is the field after the Config entry and before the keyword for the rule—in this case, the name of the rule is "ssh."

Once you've entered commands like the one shown in this example, the rule is permanently defined in the VirtualBox configuration information for the specified virtual machine, in this case "Ubuntu-Hardy." After you start the virtual machine, any host on your network can reach the SSH service on your virtual machine by using a command like `ssh host-system -p 2222`.

When setting up port forwarding in VirtualBox, you need to identify the port on the virtual machine where a given service is running (which you can do by consulting the /etc/services file), and select a port on your host system where no existing service is running. You should always use a port whose port number is greater than 1024 on your host system because port numbers less than 1024 are typically reserved for services that may actually be running on a physical system.

In the future, if you don't remember which rules you've set up, you can list any port forwarding rules that you've defined for a specific virtual machine by using a command such as `VBoxManage getextradata Ubuntu-Hardy enumerate`.

Fixing Virtual Machine Boot/Installation Problems

As mentioned earlier in this section, any user who wants to run VirtualBox must be a member of the vboxusers group. This is usually done automatically for the user who installs VirtualBox, but other users will not be able to run VirtualBox until they are added to the vboxusers group with a command like `sudo adduser USERNAME vboxusers`.

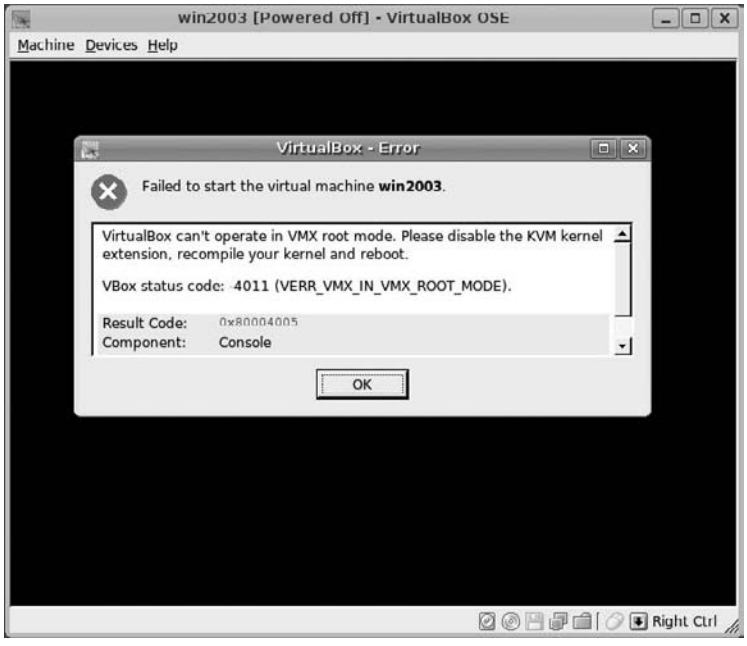
Even if the current user is a member of this group, a related problem that you may encounter is a dialog to the effect that the `/dev/vboxdrv` device was not accessible by the current user. This device should always be owned by the root user and the vboxusers group, and should be set to mode 0660 so that anyone in the vboxusers group can read/write it. See "Networking to and from Virtual Machines," later in this section, for information about a rule that you can use with your system's device creation mechanism (udev) to ensure that this device is always created with the correct permissions and user and group ownership. You can change this manually with the following commands:

```
sudo chown root.vboxusers /dev/vboxdrv
sudo chmod 0660 /dev/vboxdrv
```

Another common problem, especially if you're experimenting with multiple virtualization technologies, is that different virtualization solutions don't always play well with each other. For example, at the time of this writing, you would see an error dialog like the one shown in Figure B7-21 if you attempted to start a VirtualBox virtual machine on a system where the loadable kernel modules for the KVM virtual machine were loaded into the kernel.

FIGURE B7-21

An error dialog about conflicting virtualization solutions



Fortunately, this dialog is somewhat misleading, telling you that you have to recompile your kernel. What is this, the twentieth century? To resolve this particular problem, all you have to do is to remove the loadable kernel modules that VirtualBox is complaining about, using a command like the following:

```
sudo /sbin/rmmod kvm_intel kvm_amd kvm
```

Your system will be using only one of the `kvm_amd` and `kvm_intel` modules (because you probably have only one type of processor in your system), and will therefore complain about one of them. However, after running this command, you should be able to close the VirtualBox virtual machine window, click the Start icon in the VirtualBox window, and proceed with installing the operating system in the virtual machine. If you see a message in the virtual machine window about not being able to find any bootable media, use the Devices ➪ Mount menu command to re-mount the installation media for use by the virtual machine, and select the virtual machine window's Machine ➪ Reset command to restart the boot process.

TIP **If the loadable kernel modules for another virtualization technology are being automatically loaded when your system boots, the startup script for that technology is probably being run during system startup. Check the symbolic links in the `/etc/rc2.d` directory —you can remove the symbolic link to the "other" virtualization technology (such as KVM, in this example) to prevent its kernel modules from being loaded the next time that you boot your system.**

Installing and Using VMware Server

VMware (`www.vmware.com`) is the 800-pound gorilla of virtualization on commodity hardware, pioneering the field of commercial virtualization and providing an excellent, usable, and reliable Guest OS virtualization product for Linux and Windows desktop systems since the late 1990s. I've owned VMware Workstation since version 1.*x* in 1999 (it was a long time ago), have used their server products off and on for years, and have nothing but good things to say about their product line. Their release of VMware player and server as free software (not to be confused with open source software) was both a boon to zillions of SOHO (Small Office, Home Office), SMB (Small and Medium Business), and other acronym users, and a clever marketing move. If only more commercial software entities would see the benefits of this approach, the universe would be a better place.

NOTE **This little paean to VMware does not diminish my belief that Xen is a superior virtualization solution. Open source is "the right thing," and I believe that Xen's hypervisor is superior to the one used in VMware's ESXi product (and Microsoft's Hyper-V). A free version of VMware ESXi is now available from the VMware site, as a free version of Microsoft's is from its site. So many hypervisors, so little time …**

VMware has an extensive product line that is knee-deep in marketing language, different performance and support levels, and different permutations of management tools, and is only slightly less complex than a map of the human genome. See their Web site for more information.

In general, VMware's virtualization products fall into two groups, their Guest OS Workstation and Server products. Both support the installation and execution of unmodified x86-based operating systems, including *BSD, all Linux distributions, all flavors of Microsoft Windows, all x86 Solaris implementations, and even hacked versions of Mac OS X (which are not legal at the moment). The Workstation product supports on-demand execution of various virtual machines, although you can run multiple VMs at the same time. The free and ESX Server products feature automatic VM execution at boot time and some nice remote access and remote management utilities, although the standard free version is more of a Guest OS solution than the high-performance, bare-metal virtualization environment that ESX

provides. All VMware virtualization products feature an integrated hypervisor that provides an abstract view of your hardware that the guest operating systems can leverage without requiring any modifications. VMware's virtualization products can execute unprivileged operating system code natively, which provides intrinsic performance improvements and good opportunities for code sharing between virtual machines, where possible.

Internally, VMware's Workstation and free server products use loadable kernel modules to interface between virtual machines and your physical hardware. They also include a DHCP server, and integrated NAT and bridged networking support setup that simplifies network administration and segmentation for your virtual machines.

For more information about VMware, see:

- VMware's Web site (`www.vmware.com`): The ultimate source of information about the VMware product line and specific VMware products. The documentation for specific VMware products is available at `www.vmware.com/support/pubs/`.
- VMware communities (`http://communities.vmware.com/index.jspa`): User-driven information about VMware products, tips and tricks for using them, and much more.

The next few sections explain how to obtain, install, and configure VMware Server, and how to quickly and easily create a virtual machine with usable network support. There are plenty of books and zillions of articles about VMware setup and configuration (including its excellent product documentation), but the next few sections will help you get up and running with VMware Server and a virtual machine, so that you can see just what all of the fuss is about.

Downloading and Installing VMware Server

Although the Canonical Web pages list VMware as a partner and state that you can download VMware Server through a commercial repository, I can't seem to find it anymore. This was certainly true in previous Ubuntu releases, but perhaps things have changed. This really doesn't matter because you can easily download and install VMware Server from the VMware Web site, but the downside is that you'll have to install it manually rather than being able to install a DEB package via one of the standard package management tools that are provided with Ubuntu and Kubuntu systems.

NOTE **Part of the VMware installation process is to build the two loadable kernel modules that it uses. This means that you'll have to have a C development environment (i.e., at least the build-essential package) installed on your system in order to successfully install and use VMware Server. Installing the build-essential package, and other information about software development on Ubuntu and Kubuntu systems, was discussed in Bonus Chapter 3.**

To install and use VMware Server on your system, you'll need to get the software and at least one serial number for using it from VMware. You must register with VMware to get a serial number—you can do this at `http://register.vmware.com/content/registration.html`. After providing a variety of information, including what you plan to use VMware for, what you might be interested in buying from the VMware folks, and so on, click submit. The number of serial numbers that you requested will be generated automatically and displayed on the next page—copy and paste these into a text file for future reference. This page also contains Download Now links for the current stable version (1.0.6 at the time of this writing) or the 2.0 beta. I strongly suggest going with the version 1.*x* software to get started with VMware. You can always install a beta later if you're curious or become a VMware fan—I don't think that you want to start out by being a virtual member of the VMware QA team.

Clicking the Download Now link initially takes you to a page where you must accept the EULA (End User License Agreement) for VMware Server. After clicking Yes, you finally arrive at a page where you can download VMware Server for various platforms and in various formats. You'll want to download the Binary (.tar.gz) version. Save this file to your Ubuntu and Kubuntu system when your browser gives you that option—once the file downloads, you're ready to install VMware Server.

You must install the VMware Server from the command line, so start a GNOME terminal, xterm, or Konsole window. Once this window displays, navigate to the directory where the file that you downloaded is located, and extract its contents using a command like the following:

```
tar zxvf VMware-server-1.0.6-91891.tar.gz
```

The name of the file may differ if the VMware folks update the server after this book goes to press. This command will create a directory called `vmware-server-distrib`. Change directory to that directory, and execute the `sudo ./vmware-install.pl` command to execute this script as a privileged user and begin the installation process.

If you have an existing VMware Server installation on your system, the install script will notice that fact, and will shut down any VMware-related services that are currently running on your system and then remove the old installation before installing the new one. Don't panic! The script only removes binaries and libraries, and does not remove the configuration information in the `/etc/vmware` directory or any virtual machines that you've already created (typically located in the directory `/var/lib/vmware/Virtual Machines`).

Once any existing VMware Server installation has been purged, the VMware install script installs the new software, prompting you for a variety of information, as shown in the following example from one of my installations:

```
In which directory do you want to install the binary files?
[/usr/bin]

What is the directory that contains the init directories (rc0.d/ to rc6.d/)?
[/etc]

What is the directory that contains the init scripts?
[/etc/init.d]

In which directory do you want to install the daemon files?
[/usr/sbin]

In which directory do you want to install the library files?
[/usr/lib/vmware]

In which directory do you want to install the manual files?
[/usr/share/man]

In which directory do you want to install the documentation files?
[/usr/share/doc/vmware]
```

As you can see, I typically accept all of the defaults by pressing the Enter/Return key after each prompt. Once the software itself is installed, the installation script offers to configure it for you by running another script, `/usr/bin/vmware-config.pl`. Press Enter/Return to run this script when prompted

to do so. The first step in the configuration script is to accept yet another EULA. After paging through this, type **yes** to accept the EULA and press Enter/Return to proceed to the actual configuration process.

The configuration script will prompt you for the locations in which you want to install a few additional files, as in the following example:

```
In which directory do you want to install the mime type icons?
[/usr/share/icons]

What directory contains your desktop menu entry files? These files have a
.desktop file extension. [/usr/share/applications]

In which directory do you want to install the application's icon?
[/usr/share/pixmaps]
```

I typically accept the defaults here, too, by pressing Enter/Return in response to each prompt.

Next, the real fun begins. The VMware Server configuration script checks if it has pre-built kernel modules for your kernel (which it does not, trust me). This isn't a bad thing—the Canonical folks provide updated versions of the kernels used by Ubuntu and Kubuntu systems fairly often and you can't expect an external software package to keep up with that. You will see messages like the following during the detection process:

```
Trying to find a suitable vmmon module for your running kernel.
None of the pre-built vmmon modules for VMware Server is suitable for your
running kernel.  Do you want this program to try to build the vmmon module for
your system (you need to have a C compiler installed on your system)? [yes]
```

You will want to say **yes** (or just press Enter/Return). The install script will identify the location of the C compiler that it is using, and will prompt you for the location of the header files that are associated with the kernel that is currently running on your system. You can usually simply accept its suggestion because it is derived by examining your system. The install script will then attempt to build the `vmmon` kernel modules used for interaction between virtual machines and the Linux kernel. The output of all of this looks like the following:

```
Using compiler "/usr/bin/gcc". Use environment variable CC to override.
What is the location of the directory of C header files that match \
your running kernel? [/lib/modules/2.6.24-19-386/build/include]
 Extracting the sources of the vmmon module.
Building the vmmon module.
 Using 2.6.x kernel build system.
make: Entering directory '/tmp/vmware-config0/vmmon-only'
make -C /lib/modules/2.6.24-19-386/build/include/.. \
     SUBDIRS=$PWD SRCROOT=$PWD/. modules
make[1]: Entering directory '/usr/src/linux-headers-2.6.24-19-386'
  CC [M]  /tmp/vmware-config0/vmmon-only/linux/driver.o
  CC [M]  /tmp/vmware-config0/vmmon-only/linux/hostif.o
  CC [M]  /tmp/vmware-config0/vmmon-only/common/cpuid.o
  CC [M]  /tmp/vmware-config0/vmmon-only/common/hash.o
  CC [M]  /tmp/vmware-config0/vmmon-only/common/memtrack.o
  CC [M]  /tmp/vmware-config0/vmmon-only/common/phystrack.o
  CC [M]  /tmp/vmware-config0/vmmon-only/common/task.o
  CC [M]  /tmp/vmware-config0/vmmon-only/common/vmx86.o
```

```
  CC [M]  /tmp/vmware-config0/vmmon-only/vmcore/moduleloop.o
  LD [M]  /tmp/vmware-config0/vmmon-only/vmmon.o
  Building modules, stage 2.
  MODPOST 1 modules
  CC      /tmp/vmware-config0/vmmon-only/vmmon.mod.o
  LD [M]  /tmp/vmware-config0/vmmon-only/vmmon.ko
make[1]: Leaving directory '/usr/src/linux-headers-2.6.24-19-386'
cp -f vmmon.ko ./../vmmon.o
make: Leaving directory '/tmp/vmware-config0/vmmon-only'
The module loads perfectly in the running kernel.
```

Halfway there! Next, you'll see the following prompt:

```
Do you want networking for your virtual machines? (yes/no/help) [yes]
```

Because this isn't 1981, you'll probably want networking. Press Enter/Return to have the install script set up a bridged network that includes your system's default Ethernet device. If you have multiple Ethernet interfaces available (including existing bridges), you will be prompted to identify the one with which your VMware Server's networking support should be bridged. You will probably want to simply press Enter/Return to bridge VMware's internal networking support with your system's basic Ethernet device. This interaction looks like the following:

```
Do you want networking for your virtual machines? (yes/no/help) [yes]

Configuring a bridged network for vmnet0.

Your computer has multiple ethernet network interfaces available:
     eth0, virbr0.
Which one do you want to bridge to vmnet0? [eth0]

The following bridged networks have been defined:

. vmnet0 is bridged to eth0

Do you wish to configure another bridged network? (yes/no) [no]
```

NOTE **VMware's bridging is done internally, not through the standard Linux bridges that you create using tools such as `brctl`. You will not see VMware's bridges listed in `brctl` output—although disconcerting, this is not a problem.**

Once you've finished configuring bridged networking, the install script will prompt you for information about whether you'd like to enable NAT networking in your virtual machines. I typically answer yes (i.e., accept the default suggestions) here, because this will enable you to use either bridged or NAT networking in your virtual machines. With bridged networking, your virtual machines will get IP addresses on your primary network, and the internal bridging will enable other systems to reach your virtual machines. With NAT networking, your virtual machines will be on a private subnet with addresses allocated by VMware's internal DHCP server, but will still be able to connect to external networks, though enabling other machines to reach them will require some changes as to how routing is configured on your systems.

The prompts and messages associated with enabling NAT networking in your VMware Server will look something like the following:

```
Do you want to be able to use NAT networking in your virtual \
   machines? (yes/no) [yes]
```

```
Configuring a NAT network for vmnet8.

Do you want this program to probe for an unused private subnet? \
   (yes/no/help) [yes]

Probing for an unused private subnet (this can take some time)...
The subnet 192.168.66.0/255.255.255.0 appears to be unused.

The following NAT networks have been defined:

. vmnet8 is a NAT network on private subnet 192.168.66.0.

Do you wish to configure another NAT network? (yes/no) [no]
```

After enabling NAT networking and creating a NAT network, the install script will prompt you as to whether you want to set up host-only networking, which will only enable your virtual machines to talk to the system on which they are hosted. I never set this up, and thus usually answer "no" to the host-only networking prompt:

```
Do you want to be able to use host-only networking in your \
   virtual machines? [yes] no
```

After defining the types of networking that you want to be able to use with your virtual machines, the install script will attempt to build the loadable kernel module used by VMware Server for networking support:

```
Extracting the sources of the vmnet module.
Building the vmnet module.

Using 2.6.x kernel build system.
make: Entering directory '/tmp/vmware-config0/vmnet-only'
make -C /lib/modules/2.6.24-19-386/build/include/.. \
     SUBDIRS=$PWD SRCROOT=$PWD/. modules
make[1]: Entering directory '/usr/src/linux-headers-2.6.24-19-386'
  CC [M]  /tmp/vmware-config0/vmnet-only/driver.o
  CC [M]  /tmp/vmware-config0/vmnet-only/hub.o
  CC [M]  /tmp/vmware-config0/vmnet-only/userif.o
  CC [M]  /tmp/vmware-config0/vmnet-only/netif.o
  CC [M]  /tmp/vmware-config0/vmnet-only/bridge.o
  CC [M]  /tmp/vmware-config0/vmnet-only/procfs.o
  CC [M]  /tmp/vmware-config0/vmnet-only/smac_compat.o
  SHIPPED /tmp/vmware-config0/vmnet-only/smac_linux.x386.o
  LD [M]  /tmp/vmware-config0/vmnet-only/vmnet.o
  Building modules, stage 2.
  MODPOST 1 modules
  CC      /tmp/vmware-config0/vmnet-only/vmnet.mod.o
  LD [M]  /tmp/vmware-config0/vmnet-only/vmnet.ko
make[1]: Leaving directory '/usr/src/linux-headers-2.6.24-19-386'
cp -f vmnet.ko ./../vmnet.o
make: Leaving directory '/tmp/vmware-config0/vmnet-only'
The module loads perfectly in the running kernel.
```

Resolving Problems Building VMware Kernel Modules

Problems building any VMware kernel module are generally due to differences between the kernel include files and associated data structures that the VMware kernel module source code expects, and those used in the kernel that you are running (and therefore building against). To resolve these, you have several options, which I suggest trying in order:

- Check the VMware site for patches that provide support for your kernel version. For example, to build on the 2.6.27 kernels, you must download a patch from www.vmware.com/download/modules/modules.html.
- Download and install the VMware-Any-Any package, which is designed to enable the VMware modules to compile on any kernel. You can find this package at `http://groups.google.com/group/vmkernelnewbies/files` or `http://platan.vc.cvut.cz/ftp/pub/vmware`. The latter is the traditional location of this patch, but the former has patches for new kernel versions. You should always get the latest version of the file with a name like `vmware-any-any-update`*`XXX`*`.tar.gz`, where *`XXX`* is a version number for this package. (The version available at the time of this writing is 117d.) If you have problems using this patch on newer Linux kernels, you may want to search the Web for versions of this package for the latest kernel versions. Be careful, however—some of the sites that claim to host these packages actually resolve to porn sites, which is irritating (but somewhat amusing, depending on your sense of humor).
- Download a newer version of the VMware Server.

After downloading and uncompressing a patch file, change directory to the directory that was created and execute the command `sudo ./runme.pl`. This will patch portions of the vmware-config.pl script and associated files, and will then offer to restart the VMware Server configuration script for you, which should now work correctly.

Once the `vmnet` module is compiled and loaded into the kernel, the most configurable and critical aspects of VMware Server installation are complete. You will probably want to accept the defaults for the next few prompts, until you are prompted for your VMware Server license. You can copy and paste a license from the file in which you saved your licenses, as suggested earlier in this chapter. Once you've entered a valid license, the install script starts the services that VMware requires, and exits. This interaction with the install script looks like the following:

```
Please specify a port for remote console connections to use [902]

Configuring the VMware VmPerl Scripting API.

Building the VMware VmPerl Scripting API.

Using compiler "/usr/bin/gcc". Use environment variable CC to override.
Installing the VMware VmPerl Scripting API.

The installation of the VMware VmPerl Scripting API succeeded.
```

```
 Generating SSL Server Certificate

In which directory do you want to keep your virtual machine files?
[/var/lib/vmware/Virtual Machines]

 Please enter your 20-character serial number.

Type XXXXX-XXXXX-XXXXX-XXXXX or 'Enter' to cancel: NNNNN-NNNNN-NNNNN-NNNNN

Starting VMware services:
   Virtual machine monitor - done
   Virtual ethernet - done
   Bridged networking on /dev/vmnet0 - done
   Host-only networking on /dev/vmnet8 (background) - done
   NAT service on /dev/vmnet8 - done
The configuration of VMware Server 1.0.6 build-91891 for Linux
  for this running kernel completed successfully.
```

Congratulations—you've got VMware Server! At this point, you should execute the `vmware` command from your command line, to make sure that it starts correctly and doesn't display any warnings or other messages. Unfortunately, on Ubuntu and Kubuntu systems, the VMware Server may initially display a few error messages like the following, and then exit:

```
# vmware
/usr/lib/vmware/bin/vmware: \
  /usr/lib/vmware/lib/libgcc_s.so.1/libgcc_s.so.1: \
  version 'GCC_3.4' not found (required by /usr/lib/libcairo.so.2)
/usr/lib/vmware/bin/vmware: \
  /usr/lib/vmware/lib/libgcc_s.so.1/libgcc_s.so.1: \
  version 'GCC_4.2.0' not found (required by /usr/lib/libstdc++.so.6)
/usr/lib/vmware/bin/vmware: \
  /usr/lib/vmware/lib/libgcc_s.so.1/libgcc_s.so.1: \
  version 'GCC_3.4' not found (required by /usr/lib/libcairo.so.2)
/usr/lib/vmware/bin/vmware: \
  /usr/lib/vmware/lib/libgcc_s.so.1/libgcc_s.so.1: \
  version 'GCC_4.2.0' not found (required by /usr/lib/libstdc++.so.6)
/usr/lib/vmware/bin/vmware: \
  /usr/lib/vmware/lib/libgcc_s.so.1/libgcc_s.so.1: \
  version 'GCC_3.4' not found (required by /usr/lib/libcairo.so.2)
/usr/lib/vmware/bin/vmware: \
  /usr/lib/vmware/lib/libgcc_s.so.1/libgcc_s.so.1:
  version 'GCC_4.2.0' not found (required by /usr/lib/libstdc++.so.6)
```

These problems occur because the VMware Server tarball comes with a version of the `libgcc_s.so.1` that is incompatible with the libraries that are provided with some versions of that library for Ubuntu and Kubuntu systems. This is easy enough to resolve. Do the following:

```
pushd /usr/lib/vmware/lib/libgcc_s.so.1
sudo mv libgcc_s.so.1 libgcc_s.so.1.ORIG
sudo ln -s /lib/libgcc_s.so.1
popd
```

At this point, you should be able to start the VMware Server from your command line without seeing any error messages. In the future, you will probably want to start it from your system's menus.

TIP **Almost any kernel update on your system will prevent your system from being able to reload the VMware loadable kernel modules when you restart your system after the kernel update. You'll see a message like the following at boot time:**

```
vmware is installed, but it has not been (correctly) \
configured for this system. To re-configure it, invoke \
the following command: /usr/bin/vmware-config.pl.
```

To resolve this problem, re-run the `vmware-config.pl` script after the system has booted, as suggested. This will rebuild and load VMware's loadable kernel modules, after which VMware should load correctly (until the next kernel update, that is).

Starting VMware Server

After installing VMware Server, you can start it by either typing the `vmware` command from any command-line environment, or by selecting the Applications ➪ System Tools ➪ VMware Server Console menu command on an Ubuntu system, or the K Menu ➪ Applications ➪ System Tools ➪ VMware Server Console menu item on Kubuntu systems.

NOTE **Each time that you start the free version of VMware Server that is discussed in this chapter, you will see a message stating that a newer version is available, and offering to connect you to the VMware site to get that version. If you want to follow the VMware-related instructions in this chapter, do *not* retrieve and install the newer version. The newer version is based on VMware 2, and uses a completely different, Web-based administrative interface that I find slow, irritating, and hard to use. I would suggest using any of the other virtualization solutions discussed in this chapter before trying to deal with VMware 2's Web-based administrative interface.**

TIP **If you are trying to run VMware Server for the first time and are having problems, try starting it from a command-line environment and verifying that no error messages are being displayed. See the end of the previous section for a discussion of resolving a common type of error that occurs when starting VMware for the first time on many Ubuntu and Kubuntu systems.**

When you successfully start VMware Server, you will initially see a dialog like the one shown in Figure B7-22.

This dialog enables you to identify whether you want to connect to a VMware Server application that is running on your local machine, or to a remote VMware Server that is running on some other machine. When you're creating virtual machines on your system, regardless of who will use them, you'll want to connect to your local VMware instance, so you can simply click Connect. A dialog like the one shown in Figure B7-23 displays.

Figure B7-23 shows the VMware Server Console window. It will look like this figure only if you have not created any virtual machines because you will typically see tabs or an Inventory pane that lists all of the virtual machines that are available in this VMware Server installation. See the section on "Starting a Virtual Machine," later in the VMware portion of this chapter.

Speaking of which, proceed to the next section for information about creating a virtual machine. Once you've created a virtual machine, a subsequent section explains how to start a virtual machine, whether you're doing so to install its operating system or to actually use the virtual machine.

FIGURE B7-22

Identifying the VMware Server to connect to

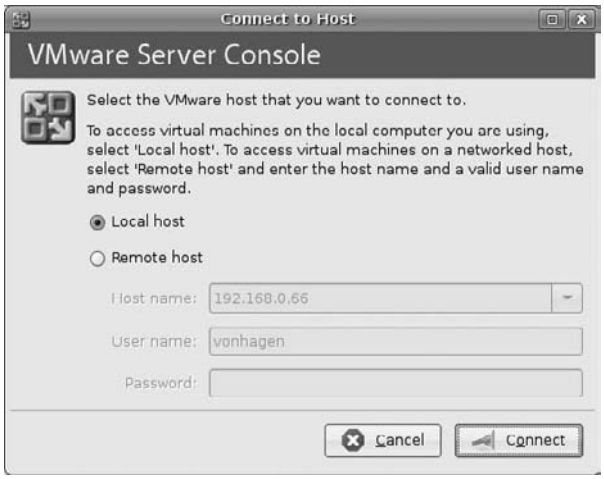


FIGURE B7-23

VMware Server's Server Console



Creating a Virtual Machine

Creating a virtual machine in VMware is fairly easy thanks to VMware's easy-to-use New Virtual Machine Wizard. To begin creating a virtual machine, either click the Create a new virtual machine button shown in Figure B7-23, or select the File ➪ New ➪ Virtual Machine menu command. This displays a generic Welcome dialog. Click Next to display a dialog like the one shown in Figure B7-24.

FIGURE B7-24

Identifying the type of virtual machine to create



This dialog enables you to specify the type of virtual machine that you want to create—a Typical one, based on predefined templates that are already available in VMware, or a Custom one, in which you will need to define characteristics such as the number of processors, who can have access to the virtual machine, the types of disk interfaces that it supports (IDE, SCSI, and manufacturer), and the type of virtual disk to create, as well as the standard questions associated with a typical virtual machine. This section focuses on creating a typical virtual machine.

To create a typical virtual machine, click Next. A dialog like the one shown in Figure B7-25 displays.

This dialog enables you to select the general type of operating system that your virtual machine will run by selecting the appropriate radio button, and then identifying the specific version that the operating system will run from the Version drop-down list. The match between version and the actual operating system that you're installing doesn't have to be exact—your selection is primarily used to select some reasonable minimum defaults for various parameters.

After selecting the general type of operating system for the virtual machine that you are creating (including "Other") and the specific or approximate version of that type of operating system, click Next to proceed. A dialog like the one shown in Figure B7-26 displays.

FIGURE B7-25

Identifying an operating system for a virtual machine

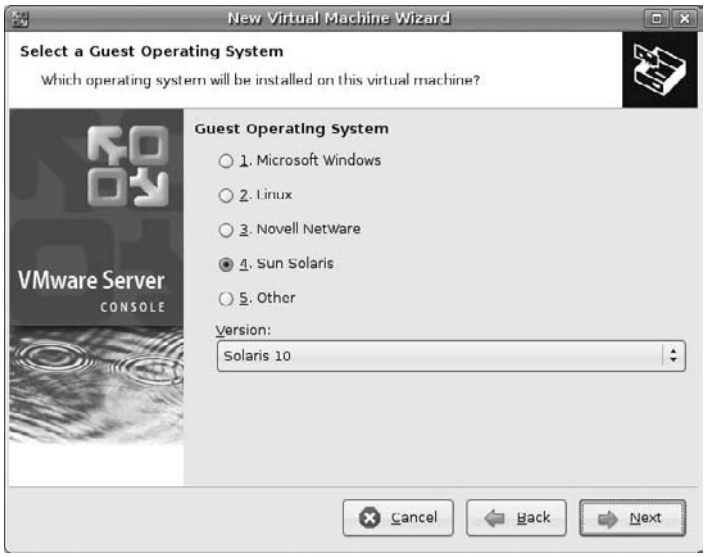
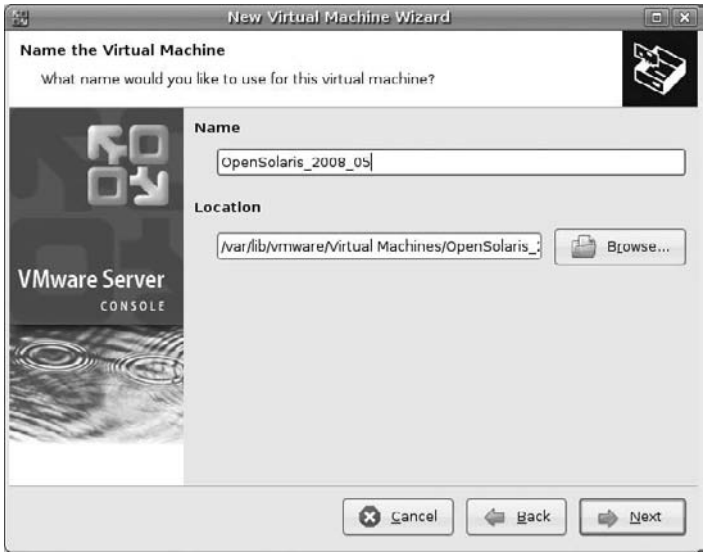


FIGURE B7-26

Naming your new virtual machine



This dialog enables you to supply a name for the new virtual machine. As you type the name, you'll notice that the name is being used not only as a label for the virtual machine, but also as the name of the directory in which the virtual machine will be stored by default. If you're paranoid about file and directory names that contain spaces, you can use whatever name you like and then edit the Location to use underscores or to concatenate words. Once you've entered the name for your virtual machine and made any optional changes to the location where it will be stored, click Next to proceed. A dialog like the one shown in Figure B7-27 displays.

FIGURE B7-27

Defining networking for your virtual machine



This dialog enables you to specify the type of networking that you want to use in your virtual machine. The following are your choices (in alphabetical order):

- **Bridged networking:** The virtual machine will appear to be on the same network as the host machine, and will obtain information from and be able to reach the same services that the host machine can. For example, if a virtual machine using bridged networking obtains its IP configuration information via DHCP, it will receive this information from any DHCP server that is running on the network to which the host is connected to, not from the DHCP server that is integrated into VMware Server.

CAUTION **As mentioned in the installation section, you do not have to set up a standard brctl bridge and define TUN/TAP devices when using VMware's bridged networking. VMware handles this internally—no fuss, no muss.**

- **Host-only networking:** The virtual machine will be connected to a private virtual network on which only the virtual machine and the host system can communicate. No external network access from the virtual machine is provided by default, though you can enable this manually by using the `route` command on the virtual machine to identify the IP address of the gateway between the private virtual network and another network (such as a corporate network or the Internet), which is typically INTERNALNETWORK.1. For example, the host-only network `172.16.80.0/255.255.240.0` would have a gateway address of `172.16.80.1`. You must then enable port forwarding on the system on which the virtual machines are running, which you can most easily do by executing the command `sudo sysctl -w net.ipv4.ip_forward=1` and then uncommenting the `net.ipv4.ip_forward = 1` command in the `/etc/sysctl.conf` file so that this happens automatically whenever you restart the host system in the future.
- **NAT networking:** The virtual machine will be on a virtual private network, and will use VMware Server's internal DHCP server to get network configuration if it uses DHCP. Virtual machines can transparently contact hosts on other networks through the internal support for NAT that is provided by VMware Server. Your host system can contact the virtual machine directly because its routing table will contain a rule for contacting the private network. External hosts can only contact the virtual machine if their routing tables identify the host system as a gateway to the private network.
- **No networking:** The virtual machine will not have network access. Nowadays, I'm not exactly sure why this would be useful. Perhaps game playing? Security through obscurity for financial data that you are maintaining in VisiCalc?

I suggest using NAT networking if you expect to access your virtual machines only from your host system, and to use bridged networking for all other cases. Although I've used VMware for a number of years, I've never found host-only networking to be very useful outside of testing scenarios where you want to guarantee that your virtual machines aren't accessing services on other machines. As the installation example in the previous section shows, I usually don't even enable host-only networking when installing VMware Server.

After selecting the type of networking that you want to use for your new virtual machine, click Next. A screen like the one shown in Figure B7-28 displays.

This screen enables you to define the characteristics of the virtual disk that will be created for and used by your virtual machine. Each virtual machine that you create has a virtual disk associated with it, in which all of the files for its operating system and any files that you create in that virtual machine are stored. These virtual disks exist as one or more files on your host system. You can create virtual disks in one of two ways, either as a fixed-size virtual disk to which all disk space is pre-allocated on the host system (by selecting the Allocate all disk space now checkbox), or as a dynamically expanding virtual disk that grows as you add more files to your virtual machine (by making sure that the Allocate all disk space now checkbox is not selected). Each of these has advantages and disadvantages:

- Pre-allocating a virtual disk provides slightly higher performance in both your virtual machine and on the host system because neither system has to allocate additional space for the disk and do the housekeeping required to accurately track the disk blocks on the host system that are actually allocated to the virtual machine. On the down side, your host system will have less storage remaining after allocating the entire disk, and your backups and any virtual machine snapshots that you create will be larger because they must back up or snapshot the entire disk image.
- Using a dynamically expanding disk results in slightly lower performance for both your host system and the virtual machine because they will have to dynamically allocate new space to the virtual disk image as it is required, and will also have to perform the housekeeping associated with that allocation. Your host system must always have enough disk space available for the maximum

size of the virtual disk image—running out of disk space when trying to expand a virtual disk image can seriously corrupt the virtual disk image, which could render the virtual disk unusable, meaning that you could lose all of the data that it contains. On the plus side, dynamically expanding images require substantially less storage on the host system, consume less space in your backups, and will require less space for any snapshots that you create of your virtual machines.

FIGURE B7-28

Defining a virtual disk for your virtual machine



Use the up and down arrows to define the size of your virtual disk if the default size isn't acceptable. This represents the maximum size of your virtual disk is you are creating a dynamically expanding virtual disk, or the amount of disk space that will immediately be allocated to your virtual disk if you use a pre-allocated disk. I suggest using a dynamically expanding disk in most cases (such as for home use) because this requires less disk space up front. However, using a dynamically expanding disk requires that you always have sufficient disk space available for the disk to grow to its maximum size. When creating virtual machines for production use, using a pre-allocated, fixed size disk is probably a better idea so that you never have to worry about running out of disk space and possibly corrupting the filesystem of a critical virtual machine.

By default, virtual disks are divided into 2GB files because many older systems couldn't support files greater than that size. I usually disable the Split disk into 2GB files checkbox because modern machines don't have a problem with larger files, and also because accessing a single disk image file provides slightly higher performance.

After specifying the characteristics of your virtual disk, click Finish to complete the virtual machine definition process. Congratulations! All you have to do now is to install your operating system of choice in the virtual machine, as explained in the beginning of the next section.

Starting a VMware Virtual Machine

Once you have defined one or more virtual machines in VMware Server, you'll see a screen like the one shown in Figure B7-22 whenever you start VMware Server, but which will also display tabs above the lower pane of the startup screen. Each tab displays the name of the virtual machine with which it is associated. If not enough space is available to display all of the tabs for the virtual machines that you have defined, you can use the arrows at the left and right ends of the tab bar to scroll left and right until you find the tab for the virtual machine that you're interested in. Selecting the tab for a specific virtual machine displays a screen like the one shown in Figure B7-29.

FIGURE B7-29

The VMware Server tab for a specific virtual machine



TIP **Immediately after you define a new virtual machine, exiting the Virtual Machine Creation Wizard causes the VMware Server console to display the tab for the virtual machine that you just defined.**

The first time you start a new virtual machine, you'll probably do so using the installation media for its operating system so that you can actually install the operating system in the virtual machine. Subsequently, you'll probably boot directly from the virtual disk for that virtual machine.

To start a virtual machine, click Power on this virtual machine. What happens next differs based on whether you are starting a virtual machine for the first time or starting a virtual machine in which you've already installed an operating system:

- If you are starting a virtual machine for the first time so that you can install an operating system in it, make sure that the installation media for the operating system that you want to install is located in the host system's CD/DVD drive. The system will boot from the CD/DVD drive by default, and will display the boot or installation screen for the operating system that you are installing. For example, Figure B7-30 shows the GRUB boot screen that displays when you boot from an OpenSolaris Live/Install CD.

FIGURE B7-30

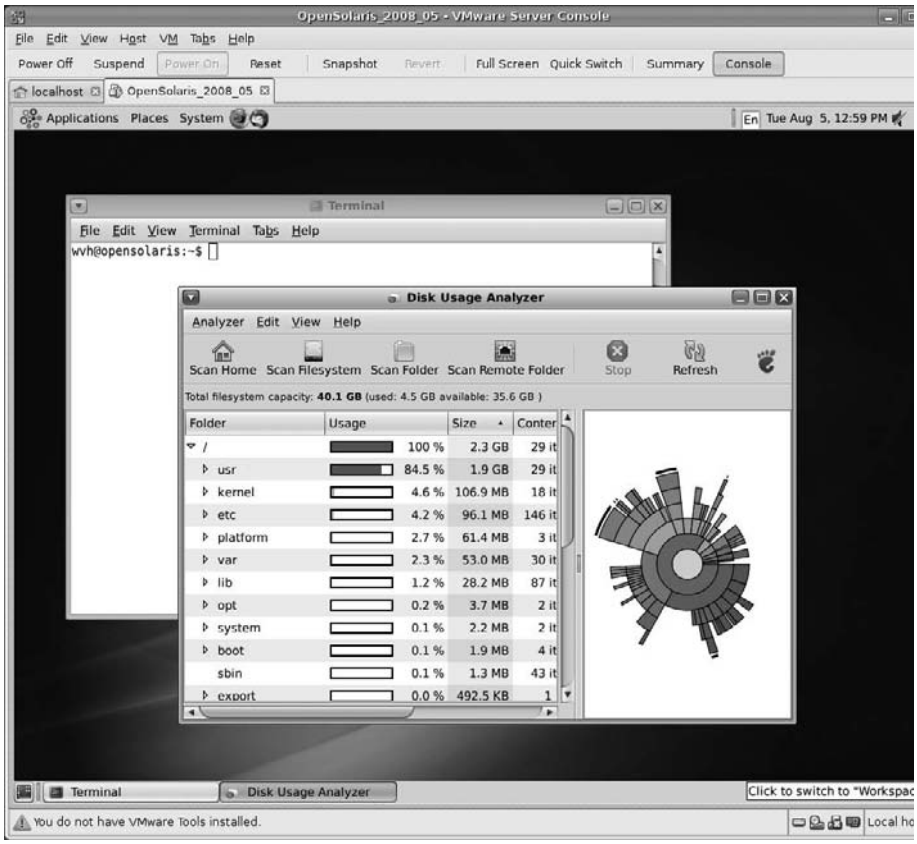
Booting a virtual machine from installation media



- At this point, the installation process is exactly like that for a physical machine on which you are installing an operating system. Once installation has completed, you should reboot the virtual machine, remove the installation media from the CD/DVD drive, and enjoy using your new virtual machine.
- If you are starting a virtual machine in which you've already installed an operating system, the virtual machine will go through the standard boot process for that type of operating system, and will eventually display a login or desktop screen. After logging in, if necessary, you can use the virtual machine just as you would use a physical machine.

FIGURE B7-31

A virtual machine running in VMware Server



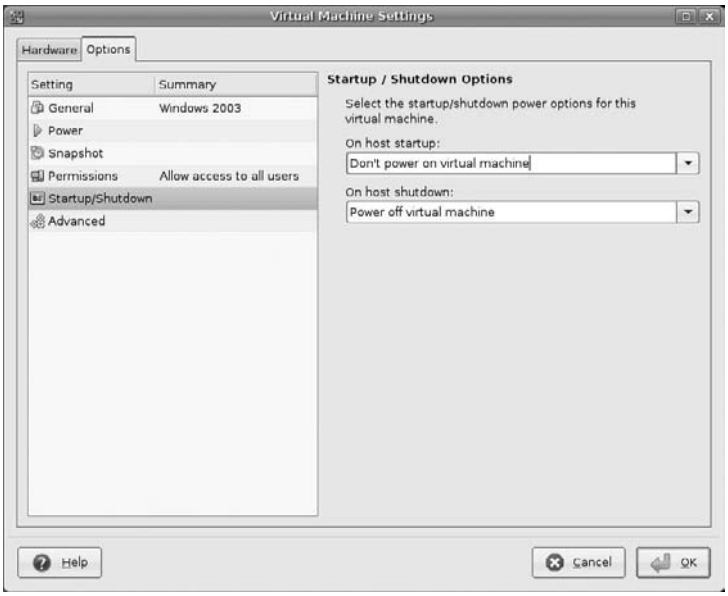
To interact with your virtual machine, simply click inside the virtual machine window. The virtual machine's cursor will begin to track the movements of your mouse, and will not be able to move outside the boundaries of the virtual machine window. This is due to VMware's Capture input capability, which makes it easy to interact with a virtual machine by sending all mouse and keyboard events directly to the virtual machine rather than the host system. This enables you to type keyboard commands such as the Ctrl+Alt+Delete required for Windows logins that might otherwise accidentally shut down your host system, and makes interacting with the virtual machine more precise because the cursor can't leave the boundary of the virtual machine window while you're in this mode. To release the mouse and cursor from your virtual machine, press the Ctrl and Alt keys at the same time. (You can change this key binding later if you want.) As a handy reminder, the keys that you must use to release keyboard capture are listed in the lower-left corner of a VMware virtual machine window when the mouse and keyboard input have been captured by that window.

If you don't always need to use your virtual machine(s), you can start both VMware and specific virtual machines whenever you need them, as described previously in this section. However, if you will always

want to be able to use your virtual machine and want to be able to access it as soon as your system boots, you can configure one or more virtual machines to start automatically when you boot your system. To do this, select the tab for a specific virtual machine and click Edit virtual machine settings. The Virtual Machine Settings dialog displays. Click the Options tab to display the general configuration options that are available for this virtual machine, and click the Startup/Shutdown settings entry in the left column to display a screen like the one shown in Figure B7-32.

FIGURE B7-32

Automating virtual machine startup in VMware Server



To configure the selected virtual machine to start automatically whenever your system restarts, select Power on virtual machine from the On host startup drop-down menu. If you do this, you should also select Shut down guest operating system from the On host shutdown drop-down menu, so that the operating system in this virtual machine shuts down cleanly whenever the host system is shut down.

Installing and Using KVM

Many Linux virtualization devotees were surprised when Linux Torvalds accepted patches that incorporated a largely unknown virtualization technology call Kernel-based Virtual Machines (KVM) from a largely unknown company (`www.qumranet.com/`) that was then in stealth mode. I certainly was. These patches were incorporated into the 2.6.20 release of the mainline Linux kernel. However surprising its appearance was, there are good precedents for an in-kernel virtualization technology such as KVM, most notably User-Mode Linux (UML), a mechanism for running an instance of Linux as a process on another

Linux system that has been supported in the mainline Linux kernel for years. Like UML, KVM supports the creation and execution of virtual Linux systems running as a separate process on a Linux system. KVM currently requires a processor that supports Intel Virtualization Technology (VT) or AMD Secure Virtual Machine (SVM/AMD-V). (The same processors required for running unmodified Guest OS systems under Xen, as discussed in the Xen section, later in this chapter.)

One of the primary reasons for the inclusion of KVM into the mainline kernel (aside from its bright, shiny, and clean implementation, of course) was the fact that KVM does not require a hypervisor, and therefore is a pure Linux solution—Linux is still the operating system that is running on the bare metal, as opposed to hypervisor-based solutions such as Xen and various VMware products. Internally, KVM works through a device node called `/dev/kvm`, which is created and managed by CPU-specific loadable kernel modules, `intel_kvm.ko` and `amd_kvm.ko`, one of which is loaded as `kvm.ko`. Opening this device creates a new virtual machine, which can then be managed through `ioctl()` calls that perform operations such as creating a virtual CPU, allocating memory, intercepting privileged calls, and so on. KVM also requires a user-space application (based on QEMU, which is discussed later in this chapter) to view and interact with each virtual machine.

KVM currently supports 32- and 64-bit Linux virtual machines, 32-bit Windows virtual machines, OpenSolaris, and most *BSD operating systems. Work on support for 64-bit Windows virtual machines is in progress.

For more information about KVM, see the following:

- **Qumranet's KVM Wiki (`http://kvm.qumranet.com/kvmwiki`):** Provides overview information and documentation/HOWTOs for using KVM on various Linux distributions.
- **Ubuntu Kernel Virtual Machine (`https://help.ubuntu.com/community/KVM`):** Ubuntu community documentation on installing and using KVM, creating virtual machines of different types, and so on.
- **KVM on Ubuntu (`www.michaeldolan.com/1030`):** One user's experience with installing and using KVM on Ubuntu 7.10. A nice hands-on reference, although it doesn't go into any detail about network setup.

Installing KVM and Related Packages

The software required to use KVM is located in a single package, the appropriately named `kvm` package. You can install the `kvm` package using your favorite package management tool, as described in Chapter 23. The `kvm` package does not include the loadable kernel modules that must be loaded in order to run a virtual machine under KVM—the `kvm.ko` and `kvm-`*`PROCESSOR`*`.ko` modules (where *`PROCESSOR`* is either `amd` or `intel`) are always provided and installed as part of the standard Ubuntu Linux kernel image. However, these kernel modules are not loaded by default on a standard Ubuntu or Kubuntu system—the `kvm` package includes the `/etc/init.d/kvm` startup and shutdown script, which loads the `kvm` module and the appropriate module for your processor.

The `kvm` package depends on a number of auxiliary packages, such as the bridge utilities (`br_utils`), IP routing (`iproute`), SDL (Simple DirectMedia Layer) video (`libsdl1.2debian`), and VGA BIOS (`vgabios`) packages, as well as some more standard system packages for sound, compression, and so on. The `kvm` package recommends installing the `qemu` and `vde2` (Virtual Distributed Ethernet) packages, and suggests some useful packages, such as the `debootstrap` package, which makes it easy to create a virtual machine on an Ubuntu or Kubuntu system (or on a Debian system, of course). Another useful

package is the `ubuntu-vm-builder` script, which simplifies creating Ubuntu virtual machines. The following would be a suitable command-line instruction to install these packages:

```
sudo aptitude -y install kvm qemu debootstrap ubuntu-vm-builder
```

Once KVM is installed, it's quite easy to create a virtual machine and associated disk, install your operating system of choice, and run the virtual machine. The next few sections explain how to manually create a filesystem image for a KVM virtual machine, and how to manually install an operating system into it.

TIP **The `ubuntu-vm-builder` script will create a disk image and install Ubuntu into it, but doesn't provide much insight into what's actually going on under the covers. If you want to use the `ubuntu-vm-builder` script to create a virtual machine (of type kvm), see the man page for the `ubuntu-vm-builder` command or its online help (`ubuntu-vm-builder --help`) for information about how to use that script.**

Creating a Disk for a KVM Virtual Machine

In order to create and use a KVM virtual machine, you must first create a virtual hard disk for use by that virtual machine. This virtual hard disk can be an existing logical volume or physical partition on your system, or a file in any filesystem that your system can access. Using a file as a virtual disk for your machine has two primary advantages:

- File-based virtual machines can easily be moved or cloned to another computer system that supports KVM by simply copying the file to the other machine.
- KVM inherits QEMU's support for a variety of clever disk formats, including the dynamically expanding QEMU-Copy-On-Write disk format. As explained in more detail later in this section, dynamically expanding disk formats enable you to create disk image files that initially only require enough space for the operating system itself, but which can expand as you create and add new files.

If you already have a free partition that you want to use with your virtual machine, you can just skip to the next section. No special formatting or allocation is required because the partition already exists and will be formatted during operating system installation, just like when installing that operating system on a physical machine. The rest of this section focuses on using KVM with file-based disk images.

You can create a virtual disk image in many different ways, my favorite of which are the following:

- By using the standard Linux/UNIX `dd` command to create a pre-allocated file of a specified size, as in the following example, which creates a 10GB file called `kvm_disk_image.img`:

  ```
  dd if=/dev/zero of=kvm_disk_image.img bs=1G count=10
  ```

- By using the `qemu-img` command to create a disk image in either raw (pre-allocated) or qcow (QEMU's copy-on-write format, which creates a dynamically expanding disk image). Because KVM and QEMU are quite intertwined, taking advantage of disk formats that were introduced by QEMU makes sense. The following is an example of using the `qemu-img` command to create a 10GB file called `kvm_disk_image.img`:

  ```
  qemu-img create -f qcow kvm_disk_image.img 10G
  ```

A qcow disk will start out very small, like any dynamically expanding disk image, but will expand as you install an operating system and create files in the virtual disk. As discussed in the VirtualBox and

VMware installation sections of this chapter, I suggest using pre-allocated disks for any critical, production virtual machines that you create (so that you don't have to worry about running out of disk space whenever the disk needs to expand), but dynamically expanding disks are fine for personal use or if you don't have much disk space and are willing to take the chance that you'll always have the disk space that you need.

Installing an Operating System in a Virtual Machine

After creating a disk image, starting the virtual machine so that you can install an operating system in your disk image is quite easy. When starting a virtual machine, you must identify the disk image, volume, or partition that it should use as its primary hard drive, any other devices (such as a CD-ROM drive) that you will need during installation, the amount of memory that you want to allocate to it, and the device to boot from if that is a drive other than the virtual machine's primary hard drive. Many other command-line options are available (of course)—because KVM depends on QEMU, see the online reference for the `qemu` command for a complete list, which you can do by executing the `man qemu` command from any xterm, GNOME Terminal, or Konsole window.

As an example of booting a virtual machine to install a new operating system, the following command will start a `kvm` process using an image file named `winxp.img` as the primary boot drive, allow access to `/dev/cdrom` on your host system as the CD-ROM drive for the virtual machine, allocate 512MB of memory to the virtual machine, and boot from the CD-ROM drive:

```
kvm -hda winxp.img -cdrom /dev/cdrom -m 512 -boot d &
```

This command starts KVM using an SDL (Simple DirectMedia Layer) window as a console for the virtual machine, as shown in Figure B7-33.

FIGURE B7-33

A KVM virtual machine being installed using an SDL console



TIP **Once you click inside an SDL KVM window, its QEMU underpinnings automatically capture the cursor, binding it to that window so that all mouse and keyboard events go directly to the virtual machine rather than to the host system. To release the cursor so that it is available to the host system and can move outside the boundaries of a KVM window, press Left Ctrl+Left Alt (hold down the left Ctrl key and press the left Alt key).**

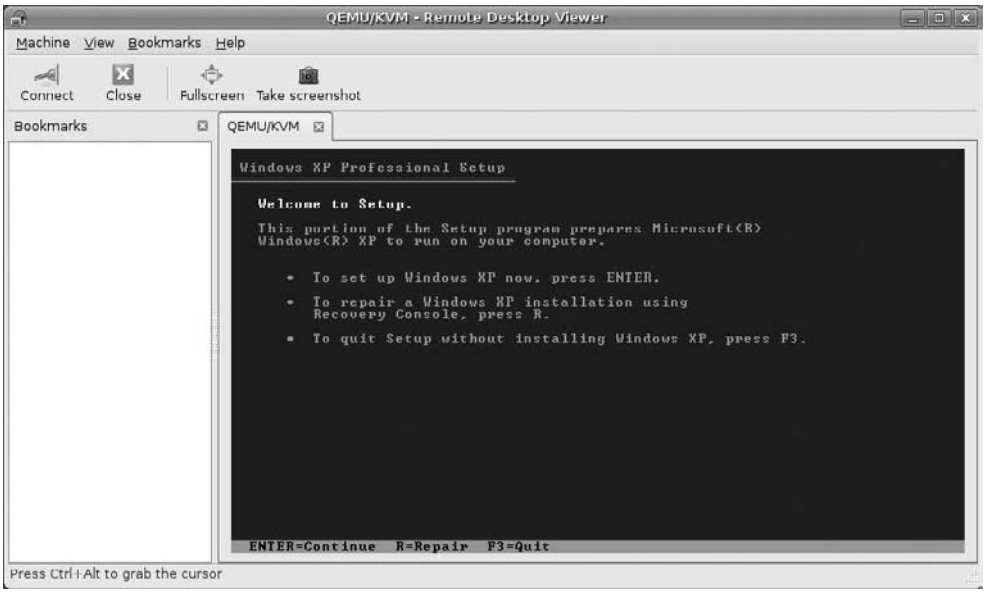
Virtual machines created via the `kvm` command can also use a VNC console rather than an SDL console, as in the following example:

```
kvm -hda winxp.img -cdrom /dev/cdrom -m 512 -boot d -vnc :1 &
```

This command starts kvm using a VNC console that you can access as `HOSTNAME:1` or `IP-ADDRESS:1` from the local system or any other system that can access the host system over the network. Figure B7-34 shows the VNC console for a kvm virtual machine in Virage, Ubuntu's latest VNC client.

FIGURE B7-34

A KVM virtual machine being installed using a VNC console



You do not need to use the `sudo` command to install or start a KVM virtual machine as long as you are a member of the kvm group (which is required to run KVM in the first place) and that you are using KVM's default networking model (i.e., the networking model that is used when you do not specify a networking model). This networking model is known as User Networking. For more information about different ways of networking KVM virtual machines, see the section "Networking for KVM," later in this chapter.

After entering a command such as either of these, the virtual machine will boot from the host system's CD-ROM drive, where you have presumably inserted the boot media for the operating system that you want to install. Once you connect to the SDL or VNC console, the installation process for that operating system is exactly the same as it would be if you were installing that operating system on a physical machine.

Starting a KVM Virtual Machine

Once you've installed an operating system in your virtual machine, you can start that virtual machine at any time using a somewhat simpler version of commands like the one shown in the previous section. The primary difference between booting a virtual machine in order to install an operating system and booting a virtual machine in order to run a virtual machine on which an operating system has already been installed is the boot device. If you simply do not specify the `-boot` option, the virtual machine will boot from its primary root filesystem, as in the following commands:

```
sudo kvm -hda winxp.img -cdrom /dev/cdrom -m 512 &
sudo kvm -hda winxp.img -cdrom /dev/cdrom -m 512 -vnc :1 &
```

Either of the previous commands will start a kvm process using an image file named `winxp.img` as the primary boot drive, allow access to `/dev/cdrom` on your host system as the CD-ROM drive for the virtual machine, and allocate 512MB of memory to the virtual machine. The first of these commands starts kvm using an SDL (Simple DirectMedia Layer) window as a console for the virtual machine. The second of these starts kvm using a VNC console that you can access as `HOSTNAME:1` or `IP-ADDRESS:1` from the local system or any other system that can access the host system over the network.

Networking for KVM

At this point, you probably won't be surprised to read that KVM's support for networking heavily leverages the networking support that is provided in QEMU. So, with yet another tip of the hat to Fabrice Bellard (Mr. QEMU), let's summarize the various networking approaches that are available in KVM:

- **User networking:** KVM/QEMU provides a DHCP service for virtual machines and serves as a packet forwarder/receiver for packets sent from the Guest OS to an external network. This provides a fairly slow network connection but requires absolutely no configuration. By default, this networking mode does not support inbound network connections to the virtual machine, but you can use the `-redir` option when starting a virtual machine to map a port on the host system to a single port on the virtual machine. User Networking is the default networking mode for KVM/QEMU if no other networking mode is specified.
- **Public or private virtual bridge via TUN/TAP:** Uses a TAP device for each virtual machine, connecting them to a virtual bridge that is created and managed using the brctl utility. If the host also belongs to the bridge (i.e., the bridge is public), this mode makes it easy for you to create virtual machines with IP addresses on the same network as the host, and therefore enables bidirectional communication between host, virtual machines, and other hosts.
- **TUN/TAP and iptables:** Associates each virtual machine with a TAP device, and uses iptables kernel packet routing rules to route traffic between an external network and each virtual machine.
- **Virtual distributed Ethernet (VDE) via TUN/TAP:** Creates a virtual Ethernet and associated switch using a single TAP device on the host. This makes it easy for you to put all of your virtual machines on a single virtual Ethernet.

This section focuses on User Networking and Public Bridging because these are the most common networking mechanisms used with KVM and QEMU. For more information about iptables, see `www.netfilter.org` or Ubuntu-specific information such as `http://ubuntuforums.org/showthread.php?t=111972`. For more information about VDE, see `http://wiki.virtualsquare.org/index.php/VDE`, or the QEMU-oriented information at `http://alien.slackbook.org/dokuwiki/doku.php?id=slackware:vde`.

NOTE **When specifying networking options for KVM or QEMU, you ordinarily use the `-net` option twice. The first option identifies the network interface card (and optionally, the device bus and VLAN) that you want to use, which defaults to an emulated PCI Realtek 8139 network interface card in KVM. The second option identifies the networking model that you want to use, which defaults to user (User Networking). The exception to the "two `-net` options" rule is when you want to indicate that no networking is supported, which is done by specifying `-net none` on the command line when you start a virtual machine.**

Using KVM with User Networking

As summarized in the previous section, the default networking mode used by KVM/QEMU is called user networking, in which KVM/QEMU serves as a packet forwarder/receiver for packets sent from the Guest OS to an external network, and also includes a DHCP server that will automatically assign an IP address to the virtual machine (on the 10.0.2.X subnet). User networking is somewhat slow, and does not support all networking protocols (no ICMP, for example). For this reason, some standard commands, such as ping, will not work outbound from the virtual machine. The biggest drawback of user networking is that it does not support inbound connections. However, if your primary reason for running a virtual machine is to enable you to occasionally run some proprietary software (some Microsoft Office component, for example), the lack of inbound connections doesn't matter, and KVM/QEMU's User networking is probably all you need. User Networking is active by default, but you can also explicitly specify it on the command line by using the `-net nic -net user` command-line options.

The User Networking support provided by KVM/QEMU doesn't prevent inbound communication to your virtual machines; it just doesn't provide full support for it. KVM and QEMU provide two mechanisms for extending the connectivity of your virtual machines when using User Networking:

- `-redir protocol:host-port:guest-ip:guest-port`: Enables you to redirect incoming connections from a port on the host system to a port on a virtual machine. `protocol` is one of tcp or udp, `host-port` is the network port on the host system that should be directed to `guest-port` on the virtual machine, and `guest-ip` is the IP address of the virtual machine if it is not the default value (10.0.2. that is assigned by the internal KVM/QEMU DHCP server). This option and associated value must be provided on the command line when you start the virtual machine. You can specify the `-redir` option multiple times on the command line to redirect multiple ports to the virtual machine.
- `-smb /full/path/to/directory`: Uses an internal KVM/QEMU SMB server to export the specified directory as a share named qemu, which the virtual machine can then mount as `\\GATEWAY-IP\qemu`. `GATEWAY-IP` is the IP address of the networking redirect inside KVM/QEMU, and is `10.0.2.2` for Windows XP and Vista virtual machines, although it may differ for other versions of Windows.

While the `-redir` option always works, I've had problems getting the `-smb` option working with different versions of KVM and QEMU. A much simpler solution is to run Samba on your Ubuntu or Kubuntu host and then configure a share named qemu in your `/etc/samba/smb.conf` file that looks something like the following:

```
[qemu]
  comment = Windows File Share for KVM/QEMU
  path = /full/path/to/directory
  browseable = yes
  read-only = no
  public = yes
```

After restarting Samba, you will then be able to mount the qemu share as `\\GATEWAY-IP\qemu` from within the virtual machine. I've had problems accessing other shares in the host environment from virtual machines, but this approach has always worked for me.

Using KVM with a Virtual Bridge

If you need to access your KVM/QEMU virtual machines from external systems and want them to be full-fledged participants on the same network as the host system on which they are running, the combination of a virtual bridge and TUN/TAP that is supported by KVM/QEMU is exactly what you need.

The steps for setting up a virtual bridge and creating a TAP device for each of your KVM/QEMU virtual machines are as follows:

1. Use your favorite package management application to install the `bridge-utils` and `uml-utilities` packages on your system. The `bridge-utils` package contains the commands that are used to create bridges, attach devices to bridges, show status information, and so on. The `uml-utilities` package contains the utilities that are used to create and manage Linux virtual Point-to-Point (TUN) and Ethernet (TAP) network devices. (The "uml" stands for User-Mode Linux, which is a mechanism for running instances of the Linux operating system as a process on a Linux machine, which was mentioned in the overview portions of this chapter.) For example, you could install these packages using aptitude from a terminal window by using the following command:

   ```
   sudo aptitude -y install bridge-utils uml-utilities
   ```

2. Modify your `/etc/network/interfaces` file to add an entry for a bridge and connect the network device that is associated with your system's physical network interface to it. You should also comment out any existing entry for that network device. The following is a sample entry for a bridge named `br0` to which the network device `eth0` is attached:

   ```
   iface br0 inet dhcp
     bridge_ports eth0
     bridge_fd 0
     bridge_maxwait 0
     metric 0
   ```

 This entry says that the `br0` is an IP bridge (because of the `bridge_ports` entry) that should get its configuration via DHCP, that no delay should be imposed when forwarding packets across the bridge (the `bridge_fd` entry), that the bridge should immediately enter forwarding mode as soon as it comes up (the `bridge_maxwait` entry), and that this interface should be considered a high-priority interface that can immediately reach a default gateway (the metric entry). Remember to comment out any existing entries for your network device (`eth0`, in the example used in this section)!

TIP **Unfortunately, Ubuntu and Kubuntu systems don't provide a man page that describes the settings that you can use to define bridge settings in an `/etc/network/interfaces` file. For more information about the settings I've used (and the many others that I did not), see the man page for the `bridge-utils-interface` additions, which you can find on the Net at locations such as `http://web.iesrodeira.com/cgi-bin/man/man2html?bridge-utils-interfaces+5`.**

3. Modify the rule file that is used when devices associated with VirtualBox are created during your system's boot process, adding a rule that specifies that the TUN/TAP virtual network device `/dev/net/tun` will be writable by anyone in the kvm group. You will want to add the following entry to the end of the `/etc/udev/rules.d/45-kvm.rules` file. Enter this on a single line, with no `\`:

```
KERNEL=="tun", NAME="net/%k", OWNER="root", \
     GROUP="kvm", MODE="0660"
```

4. If you're using an Ubuntu system, disable the GNOME network-manager. I'm a big fan of this tool in standard networking situations, but it just doesn't understand bridging. See the "Networking and VirtualBox" section, earlier in this chapter, for a more verbose discussion of potential problems—no sense in simply repeating it here. There are many ways to disable the network manager, the easiest of which is to immediately cause it and its dispatcher subprocess to exit when they are first started. To do this, execute the following two commands:

```
sudo echo "exit" > /etc/default/NetworkManager
sudo echo "exit" > /etc/default/NetworkManagerDispatcher
```

 The next time you start your system (which we'll do in the next step), no more network manager! You can also simply remove the network-manager package, but who knows if you may want to use it eventually—it might be fixed to deal with bridging, after all. This approach to disabling it is also easy to undo, by deleting these two files and then restarting the dbus command script using the command `/etc/init.d/dbus restart`.

5. Reboot your system. When your system comes back up, you can use the `brctl show` command to verify that a bridge has been created and that the right physical and virtual network devices are attached to it. The output associated with the example that I'm using looks like the following:

```
$ brctl show
bridge name bridge id           STP enabled interfaces
br0         8000.0019d1239f01   no           eth0
```

 You can also verify that routing is set up correctly by executing the `route -n` command—you should see only a few entries, all related to the bridge.

6. Start your virtual machines using a command like the following:

```
sudo kvm -hda winxp.img -cdrom /dev/cdrom -m 512 \
   -net nic -net tap
```

After entering your password, kvm will start your virtual machines, and the boot process will create a tap device for the virtual machine and execute the `/etc/kvm/kvm-ifup` script to attach it to the bridge. If your virtual machine gets its network configuration information from DHCP, it will get that information from a DHCP server on the same network as your host, not from KVM/QEMU's internal DHCP server.

Installing and Using Xen

Xen is the best known and most widely used hypervisor-based virtualization technology in use today, supporting both paravirtualization and hardware virtualization. VMware's ESXi and ESX Server products use a hypervisor of sorts, but do not offer the power and flexibility of Xen. Microsoft's Hyper-V hypervisor is in its infancy and is outside the traditional areas of expertise for Microsoft (whatever those are). Free versions of the VMware and Microsoft hypervisor offerings are available, but Xen is an open source virtualization technology that is released under the GPL. In my opinion, if you are serious about deploying a powerful, enterprise-ready virtualization technology, Xen is the right way to go.

NOTE **The Ubuntu folks seem to disagree with my advocacy of Xen as "the" enterprise virtualization technology. In an article in the Ubuntu Weekly news in Spring of 2008 (`https://wiki.ubuntu.com/UbuntuWeeklyNewsletter/Issue76`), the leader of the virtualization team for the Ubuntu Server project explained why the team was focusing on KVM as opposed to Xen, stating "We've chosen to settle on KVM as our main virtualization focus." Bzzzt! I think that this is just plain wrong. No one can argue that KVM is easier and is already in the kernel, but it's really more of a desktop, virtual machine technology than an enterprise VM technology, no matter how cool it is. KVM certainly does not provide the granular resource management, dynamic resource reallocation, trusted platform module support, and inter-machine migration capabilities of Xen, to name just a few missing features. KVM is also forever tied to the Linux kernel—which is fine for me, but may be a sore spot for the thousands of *BSD and OpenSolaris users out there, as well as the remaining 18 Solaris users.**

The Ubuntu repositories already provide Xen-enabled kernels and related tools, and I think it's just a matter of time before Xen becomes the favored virtualization technology on Ubuntu, just as it already is on Fedora, Red Hat, and SUSE systems.

Overview of Xen

Xen has its roots in the XenoServers project (`www.xenoservers.net`) at the University of Cambridge in England. The goal of the XenoServers project is to develop a powerful, flexible infrastructure for global distributed computing, which requires the ability to enable single machines to run different, or multiple, operating system instances and their associated applications in isolated, protected environments. The Xen Virtual Machine Monitor (`www.cl.cam.ac.uk/research/srg/netos/xen/`) was originally developed to support this capability for the XenoServers project, and was released as open source software under the GPL. The first public release of Xen (1.0) was in October of 2003. In 2005, the developers of Xen at the University of Cambridge founded XenSource, Inc. (`www.xensource.com`) to further the open source Xen community, and to sell and support enterprise solutions based on Xen technology. XenSource was acquired by Citrix in 2007.

Today, Xen is in use in thousands of commercial and research environments, and is commercially supported by Citrix, Virtual Iron, and a number of other firms. Xen is used commercially as a powerful, cost-effective solution to control and manage in-house IT infrastructure and related costs, and as an enabling technology for a variety of ISPs and similar service providers, where Xen simplifies providing privileged access to per-customer virtual servers.

Common Xen-related terminology such as a hypervisor and paravirtualization were introduced earlier in this chapter, in the section "Types of Virtual Machines." Another common term you'll encounter in Xen literature is the idea of a "Xen Domain," which is a specific instance of a Xen virtual machine running on a specific physical piece of hardware. Xen supports two basic types of domains, with different uses and capabilities.

Xen manages access to memory and hardware resources through a combination of its hypervisor, a specially privileged Xen-modified kernel, and a root filesystem that includes user-space utilities that enable you to manage, monitor, and administer all other Xen virtual machines running on a specific physical piece of hardware. This specially privileged Xen kernel—its root filesystem, and the Xen utilities that it contains—is known as "domain0" (or "Domain-O," or "dom0" to its friends). The domain0 kernel and the Linux system that runs on it host a Xen daemon (`/usr/sbin/xend`, written in the Python programming/scripting language) and administrative software that uses the Xen daemon to communicate with the hypervisor. Out of the box, Xen provides a command-line administrative application, known as xm (for Xen Management), in which you can initiate commands to create, connect to, and shut down other domains, and control the amount of memory that is allocated to a domain, associated scheduling parameters, and so on.

The domains started by a domain0 system are known as guest domains, unprivileged domains, or simply domainU domains. When a domain0 system creates, connects to, modifies, or terminates any domainU domain, it uses a configuration file that provides detailed information about the guest domain, such as the physical memory that is initially allocated to that domain, the access that the guest domain had to physical disks and network devices on the domain0 system, how networking is configured on the guest domain, any graphical support that you can use to connect to the console of the guest domain, and so on. Xen domain configuration files are discussed in the section "Creating Xen DomainU Configuration Files," later in this chapter.

NOTE **When Xen was first released, the kernels used by domain0 and domainU systems had to be compiled with different kernel configuration options in order to differentiate between them. This is no longer the case: a standard Xen kernel can detect whether it is being loaded by the hypervisor or through a domain0 system, and behaves accordingly.**

Beyond whether your CPU supports hardware virtualization, the ultimate arbiter of the general types of operating systems that you can run on a specific Xen domain0 host is the file `/sys/hypervisor/properties/capabilities`, which you will find on any Xen domain0 system. Entries in the `/sys` filesystem are dynamically generated files that provide hardware and system-level information about the state of a Linux system. You can use the standard `cat` command to view the contents of this file, as in the following example from a 64-bit Intel-VT system:

```
# cat /sys/hypervisor/properties/capabilities
xen-3.0-x86_64 xen-3.0-x86_32p hvm-3.0-x86_32 \
               hvm-3.0-x86_32p hvm-3.0-x86_64
```

In this case, the contents of the file say that this system can run paravirtualized 64-bit domains (xen-3.0-x86_64), paravirtualized 32-bit domains (xen-3.0-x86_32p), unmodified 32-bit operating systems as hardware virtual machines (hvm-3.0-x86_32), paravirtualized 32-bit operating systems as hardware virtual machines (hvm-3.0-x86_32p), and unmodified 64-bit operating systems as hardware virtual machines (hvm-3.0-x86_64).

A good general rule is that systems with hardware support for virtualization that are running Xen are generally backward-compatible. In other words, a Xen system with more capabilities can generally run a Linux kernel with the same or lesser capabilities. A 32-bit Xen kernel with PAE support can run a Linux kernel without PAE in a hardware virtual machine. A 64-bit Xen system can run any 32-bit Linux kernel in a hardware virtual machine. And so on.

NOTE **One other advantage of Xen over other virtualization solutions is that its network setup is automatically integrated into the startup process for both the domain0 system and each domainU system. By default, Xen uses bridging, but creating the bridge, adding the domain0 Ethernet interface(s) to it, and adding domainU interfaces to it at startup time are all handled automagically by scripts that are provided with Xen and internally executed as needed. No fuss, no muss! Xen also supports NAT and routed networking, but exploring those in detail is outside of the scope of this section.**

Installing Xen

Packages to install the Xen hypervisor, a Xen-aware kernel, and related libraries and utilities are provided in the Ubuntu repositories and can be installed on any Ubuntu or Kubuntu desktop system. The repositories provide two meta-packages to simplify installing all of the related packages that you'll need: `ubuntu-xen-desktop` (which currently has some dependency problems) and `ubuntu-xen-server`, which is really all that you'll need. The only difference between the two is that the `ubuntu-xen-desktop` package installs the xenman tools, which I never use, so I suggest installing the `ubuntu-xen-server` package.

To install Xen on an Ubuntu or Kubuntu desktop or server system, use your favorite package management software (Synaptic on Ubuntu, Adept on Kubuntu, or the command-line tools anywhere), as described in Chapter 23. During the installation process, your GRUB configuration file will be updated to boot the Xen hypervisor and associated domain0 kernel by default, and the S20xend and S21xendomains symbolic links will be created in the `/etc/rc2.d` directory, pointing to the two Xen startup scripts. The xend script starts the xen daemon, and the xendomains script enables you to automatically start Xen domains each time you boot your domain0 system.

After installing the `ubuntu-xen-server` package, reboot your system. The Xen hypervisor and kernel will load, booting your normal root filesystem and starting the Xen daemon. You're now ready to define your first Xen virtual machine, as described in the next section.

TIP **If the Xen daemon doesn't start after rebooting your machine, double-check that the appropriate startup script links have been created in the `/etc/rc2.d` directory. If the S20xend and S21xendomains links aren't present, you can create them with the following commands:**

```
cd /etc/rc2.d
sudo ln -s ../init.d/xend S20xend
sudo ln -s ../init.d/xendomains S21xendomains
```

Creating Xen DomainU Configuration Files

Each Xen virtual machine is defined in a configuration file that tells the Xen Management utility, `xm`, how to allocate resources, where to find the disk image (for hardware virtual machines) or kernel and root filesystem (for paravirtualized virtual machines), and so on. This section explains the basic format of a Xen configuration file for both paravirtualized systems and hardware virtual machines. Subsequent sections on creating each type of Xen virtual machine will provide a specific example of a Xen configuration file for the system that is created in each section.

Xen virtual machine configuration files provide many options for fine-tuning memory use, networking, and your interaction with the VM. Luckily, the majority of these are optional because Xen will substitute intelligent default values for any unspecified configuration parameters.

Configuration Files for Paravirtualized VMs

A minimal template configuration file for a paravirtualized Xen virtual machine is as follows:

```
kernel = "/full/path/to/kernel"
ramdisk = "/full/path/to/initrd"
memory = 256
name = "VM-NAME"
vif = [ '' ]
dhcp = "dhcp"
disk = ['tap:aio:/full/path/to/RFS/image,xvda1,w',
        'tap:aio:/full/path/to/swap/file.swap,xvda2,w']
root = "/dev/xvda1 ro"
extra = 'xencons=tty'
```

The entries in this file have the following meaning:

- `kernel`: Defines the location of the kernel that the virtual machine should boot.
- `ramdisk`: (optional) Defines the location of the initial RAMdisk that the virtual machine should load during the boot process.
- `memory`: The amount of memory to be allocated to the domainU, in megabytes.
- `name`: The name given to the domainU VM.
- `vif`: Identifies the characteristics of the virtual network interface. In this case, its presence tells Xen that a virtual network interface should be created, but its lack of any specific directives tells Xen to use intelligent default values for things like the virtual interface name, its MAC address, the network bridge to connect to, and so on.
- `dhcp`: Tells the VM to obtain its Ethernet address via DHCP.
- `disk`: Identifies the type, location, device, and mode of each disk object that the VM will use by default. See the "Specifying Xen Disks and Network Interfaces" sidebar for more information.
- `root`: Identifies the root device to the domainU kernel and provides initial mount arguments (just as found on the kernel line in the GRUB boot stanza for an unmodified kernel).
- `extra`: Enables you to specify additional arguments to be passed to the kernel when it boots. In this case, the `'xencons=tty'` argument indicates that the paravirtualized Xen virtual machine should use a TTY device as its console.

The key changes that you'll have to make to this template file are to supply the full pathnames for the kernel, initial ramdisk, and the disk and swap file images. You should also change the name entry to provide a name for your virtual machine that better reflects its type and/or function.

Many other options can be specified in the configuration file for a Xen hardware virtual machine. For complete information about these options, see the Xen documentation, which you can find in `/usr/share/doc/xen` after the `xen-docs` package for the release of Xen that you are using has been installed on your system (xen-docs-3.2, at the time of this writing).

Specifying Xen Disks and Network Interfaces

The Xen configuration file entries for disks/partitions and network interfaces are more complex than the other entries in a Xen VM configuration file:

- `disk`: The disk entry is a comma-separated list of triplets of the form `DEV,VDEV,MODE[,DOM]`, enclosed within square brackets, that defines the disk device(s) for a domain. The first value, `DEV`, consists of a `device-type:device-identifier` pair. Default device types are `file` (deprecated, referring to a partition image mounted as a loopback device), `tap:aio` (referring to the blktap driver, and the replacement for `file:` entries in earlier releases of Xen), and `phy` (a physical device). `DEV` is exported to the target domain as `VDEV`. The `MODE` is one of `r` (read-only), `w` (read-write), or `w!` (forced read-write). The `w!` should be used with extreme caution because it forces Xen to be able to write to the device, and should rarely be necessary. An optional fourth value, `DOM`, defines a specific backend driver domain to use for the disk. One important difference between disk entries in the configuration files for paravirtualized and hardware virtual machines is that disk entries in the former represent disk partitions, while disk entries in the latter represent entire disks.
- `vif`: the `vif` entry is a comma-separated list of entries of the form `type=TYPE,mac=MAC,bridge=BRIDGE,ip=IPADDR,script=SCRIPT,backend=DOM,vifname=NAME` that identify the characteristics of the virtual network interface(s) for a domain. If the type option is not specified, the default is `netfront`, which is a paravirtualized front-end device driver for network devices under Xen. If the `mac` option is not specified, a random MAC address of the form `00:16:3E:XX:XX:XX` is generated. If the bridge option is not specified, the first bridge found is used, which is `xenbr0` by default. If the `script` option is not specified, the default script specified in the Xen daemon's configuration file is used to configure the new virtual network interface. If the `backend` option is not specified, the default backend network driver domain is used. The `vifname` option enables you to name the virtual network interface that is associated with this domain; if none is specified, the virtual interface will have a name of the form `vifD.N`, where `D` is the domain ID and `N` is the interface ID.

The disk entry in a Xen configuration file typically identifies at least the root and swap partitions for a paravirtualized Xen VM, and typically identifies a CD/DVD drive and the disk containing those partitions for hardware virtual machines.

The most common entry that is specified in the `vif` entry in a Xen configuration file is the name of the bridge to which the virtual machine's network interface should be attached. You may also want to specify a MAC address for use with DHCP or network tracking, or a custom virtual interface name (`vifname`) to simplify network monitoring and accounting. As mentioned in the introductory material for this section, Xen's integrated bridge and virtual interface setup provides a significant administrative advantage over other virtualization solutions, but you may still want to specify `vif` options to fine-tune your domainU systems to suit the administrative requirements of your enterprise.

Configuration Files for Hardware VMs

A template Xen domainU configuration file for a hardware virtual machine looks like the following:

```
import os, re
arch = os.uname()[4]
```

```
if re.search('64', arch):
    arch_libdir = 'lib64'
else:
    arch_libdir = 'lib'
kernel = "/usr/lib/xen/boot/hvmloader"
builder='hvm'
memory = 256
name = "HVM-NAME"
vif = [ 'ioemu' ]
disk = [ 'tap:aio:/full/path/to/complete/disk/image,xvda,w',
         'phy:/dev/hda,hdc:cdrom,r' ]
device_model = '/usr/' + arch_libdir + '/xen/bin/qemu-dm'
boot="dc"
serial='pty'
usb=1
usbdevice='tablet'
sdl=1
```

The entries in this sample HVM configuration file have the following meanings:

- The first section is a block of Python code that imports two modules, identifies the architecture on which Xen is running, and sets the library directory used by Xen (standard 32-bit or 64-bit) appropriately.
- `kernel`: Defines the initial code that should be loaded when booting the virtual machine, in this case a hardware virtual machine loader that is used to emulate a bare x86 BIOS and run the target operating system's bootloader from the boot media.
- `builder`: Defines the function that should be used to construct the domainU system.
- `memory`: The amount of memory initially allocated to the HVM.
- `vif`: Specifies that this network device will use QEMU's ioemu emulation device for the network interface rather than Xen's virtual frontend network driver. See the "Specifying Xen Disks and Network Interfaces" sidebar for more information about virtual network interface specifications.
- `disk`: Identifies the storage that will be associated with each local storage device that this HVM has access to. Note that filesystem image file, disk, or other device references in an HVM configuration file identify storage that the HVM will see as an entire disk, not as a single partition. Each device statement, enclosed within single parentheses, identifies a single local storage device using the standard domain0 name, domainU name, and read/write information nomenclature. See the "Specifying Xen Disks and Network Interfaces" sidebar for more information about disk specifications.
- `device_model`: The real-mode instruction emulator used by HVMs, which is the full path to the QEMU device model binary.
- `boot`: Identifies the order in which the HVM tries to boot from specified devices, where `c` represents a virtual hard disk and `d` represents a CD-ROM drive. In this case, the virtual machine will first try to boot from the CD-ROM drive, then from the specified disk.
- `serial`: Redirects serial console output from the HVM to a pseudo-TTY, making it possible to collect console boot information by connecting to `/dev/pty/N` (where `N` is the domain ID), which you can connect to using the `xm` command or a terminal emulator such as minicom.

- `usb`: Enables you to subsequently attach USB devices to an HVM by using commands such as `xm block-attach` (for USB block storage devices) and `xm network-attach` (for USB network devices).
- `usbdevice`: Identifies how a USB mouse will track within the HVM. The `tablet` setting forces absolute movement tracking with the HVM window. The other possible value, `mouse`, tracks relative movement and uses standard techniques such as acceleration.
- `sdl`: Tells Xen to start the virtual machine by opening a Simple DirectMedia Layer (SDL) window for use as a console. If you do not specify the `sdl` option, the hardware virtual machine will provide a default console that you can access via VNC.

The key change that you'll have to make to this file is to supply the full pathnames for the disk image, physical disk, or disk partition that you want to use with this virtual machine. You should also change the name entry to provide a name for your virtual machine that better reflects its type and/or function.

Many other options can be specified in the configuration file for a Xen hardware virtual machine. For complete information about these options, see the Xen documentation, which you can find in `/usr/share/doc/xen` after the xen-docs package for the release of Xen that has been installed on your system (xen-docs-3.2, at the time of this writing).

Building and Booting a Paravirtualized DomainU

A paravirtualized domainU system typically has a few basic components:

- A Xen-aware domainU kernel and associated loadable kernel modules.
- An initial RAM disk that the domainU kernel can use to pre-load any drivers, initialize services, and so on. This is technically optional, but unless you build a kernel yourself, most Xen-aware kernels (including those in the Ubuntu repositories) expect an initial RAM disk.
- A root filesystem that contains the startup scripts, applications, and associated libraries that the domainU requires.
- A swap file or partition that can be used for virtual memory (optional).
- A configuration file that ties all of the preceding things together, identifies memory requirements, specifies networking options, and so on.

It's easy enough to put all of these things together manually, but luckily, that isn't necessary. The xen-tools package includes a script called `xen-create-image` that will use a kernel and initial RAM disk that you specify, create a root filesystem (and swap file), and also generate a configuration file that ties all of them together.

CAUTION **If you are creating filesystem image files for the root filesystem of your virtual machine, you will need to make sure the loadable kernel module for loopback devices, known as `loop`, is loaded into your kernel before executing the `xen-create-image` script. You can do this by executing the command `sudo modprobe loop`. You should make sure that this is also always loaded at boot time, which you can do by adding the entry `loop max_loop=255` to the file `/etc/modules` on an Ubuntu or Kubuntu system. The `max_loop` variable defines the maximum number of loopback devices that can be created on a single system, and 255 is its maximum value.**

A typical example of executing the `xen-create-image` script from the command line is as follows:

```
sudo xen-create-image --hostname=ubuntu1.vonhagen.org \
     --dir /home/xen --dist hardy --dhcp \
     --mirror=http://archive.ubuntu.com/ubuntu \
     --kernel=/boot/vmlinuz-2.6.18-xen \
     --initrd=/boot/initrd-2.6.18-xen.img
```

These options have the following meanings:

- `--dhcp`: Specifies that the virtual machine will get its IP address via DHCP. The alternative is to use the `--gateway` and `--ip` options to assign a static IP address.
- `--dir`: Specifies the name of the directory under which filesystem and swap images will be created. The alternative is to specify a logical volume group or EVMS container in which to store the images.
- `--dist`: Specifies the Ubuntu release that you want to use as the basis of your virtual machine. Your virtual machine will contain the default packages and the versions of those packages that are associated with the specified release.
- `--hostname`: Identifies the name of the virtual machine, which is used in several places. It will appear as:
 - The hostname of the virtual machine when it boots.
 - The basename of the configuration file for the virtual machine that the `xen-create-image` script created in `/etc/xen`. The name of this file will be `HOSTNAME.cfg`.
 - If you are creating image files, the hostname will also be used as the name of the directory in which the root filesystem and swap images will be created. This directory will be created in the domains subdirectory of the directory that you specify using the `--dir` option.
- `--initrd`: Specifies the full pathname of the initial RAM disk that you want to use with the kernel that you specified for the virtual machine using the `--kernel` option.
- `--kernel`: Specifies the full pathname of the kernel that you want to use for the virtual machine.
- `--mirror`: Specifies the URL of a server and directory that contains the packages for the release that you specified using the `--dist` option.

These are only a subset of the available options. For a complete list, see the man page for the xen-create-image file or view its usage message by executing the `xen-create-image --help` command.

After executing the `xen-image-create` command with the arguments that are appropriate for your system, you'll see a variety of messages as the configuration file, root filesystem, and swap file for your virtual machine are created.

TIP **If you're creating multiple virtual machines that are essentially the same, you can modify the `/etc/xen-tools/xen-tools.conf` configuration file that is used by this script. Modifying the configuration file enables you to set the default values that you want to use for various parameters, so that the only parameter you'll actually have to specify is the name of the virtual machine.**

Unfortunately, Xen is a moving target, and the version of the xen-create-image file that is provided in the current xen-tools package (xen-tools-3.2, at the time of this writing) doesn't take advantage of the latest Xen developments and uses a few legacy references. After running the `xen-create-image` script, you may need to use a text editor to edit the Xen configuration file that it generated, and make two updates:

- Change every example of `file:` to `tap:aio:`.
- Change the extra entry to `2 xencons=tty`.

Now you're ready to go!

To boot your new virtual machine, simply execute a command like the following (using the virtual machine created earlier in this section as an example):

```
sudo xm create /home/xen/domains/ubuntu1.vonhagen.org
```

You will see a message to the effect that the domain is being started using the specified command, and the command prompt will redisplay. You can attach to its console (and log in, once the boot process completes) by using a command like the following:

```
sudo xm connect 6
```

In this example, 6 is the domain identifier (ID) that is displayed for the new domain when you use the `xm list` command, as in the following sample command and output:

```
$ sudo xm list
Name                        ID   Mem VCPUs      State   Time(s)
Domain-0                     0  1748     2     r-----    462.7
ubuntu.vonhagen.org          6   128     1     -b----      2.5
```

While re-attaching to the console of a paravirtualized system that has already booted works fine, I prefer to automatically create a console window for each virtual machine because this makes it easy to identify any problems that arise during the boot process. To do this, execute a command like the following to start a virtual machine:

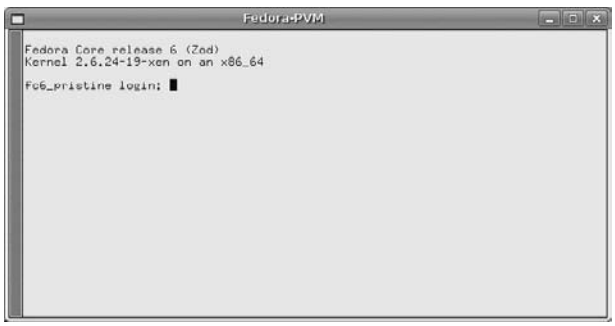
```
sudo xterm -e xm create /home/xen/domains/ubuntu1.vonhagen.org -c
```

After prompting you for your password, this command starts an xterm window in which the virtual machine is created, and which is also used as the console (`-c`) for the virtual machine. Once you've supplied your password, you can temporarily suspend the terminal process using the Ctrl+Z key combination, and then placing it in the background by using the `bg` command. Figure B7-35 shows an xterm that is serving as the console for a paravirtualized system, an old Fedora Core system in this case.

Once you're done with the virtual machine, you can connect to its console and shut it down using standard commands such as `shutdown -h now`.

FIGURE B7-35

A paravirtualized Xen domainU console window



GUIs in Paravirtualized Systems

Recent advances in Xen have enabled paravirtualized systems to actually run X servers and display console windows, but this is a huge waste of cycles in my opinion. The best way to provide a graphical user interface for a paravirtualized system with minimal impact on its performance is to install a VNC server on the paravirtualized system and start that server at boot time. I typically install the tightvncserver package, but the vnc4server package is an excellent alternative. You will also want to install any X Window system applications that you want to use, such as a window manager (hint: fluxbox), an X11 terminal emulator (hint: xterm), and any of your other favorite X11 applications.

Once started, the VNC server will export a graphical desktop (based on the commands listed in `~/.vnc/xstartup`) which you can access from any VNC client application. The following is a sample xstartup file that loads personalized X Window system resources, sets the background to gray, starts two xterms, and then starts the Fluxbox window manager:

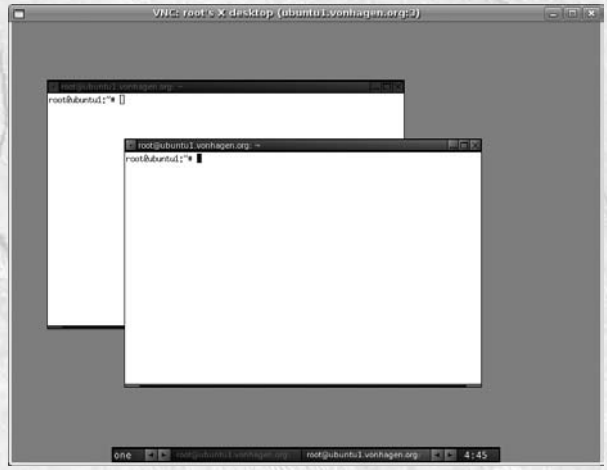
```
#!/bin/sh
xrdb $HOME/.Xresources
xsetroot -solid grey
xterm &
xterm &
fluxbox &
```

Figure B7-36 shows an X11 desktop from the paravirtualized system that was created using the xen-create-image script, as described in the previous section. The VNC server on this system uses the sample xstartup file that was just discussed.

Users can then connect to this graphical desktop using any VNC client. Access to this graphical desktop can be password protected by executing the `vncpasswd` command on the virtual machine.

FIGURE B7-36

A graphical desktop for a paravirtualized domainU system



TIP **If you are going to be creating Ubuntu-based virtual machines, Ubuntu provides a script called `ubuntu-vm-builder` to simplify the process. This script is not installed by default, but is located in the `ubuntu-vm-builder` package. Unfortunately, at the time of this writing, this script doesn't understand how to build Xen virtual machines. Once the Ubuntu folks wise up to the pre-eminence of Xen, perhaps they'll fix the script. Or you can! Here's your opportunity to contribute to the world of open source…**

Creating a Hardware Virtual Machine DomainU

If you've read other portions of this chapter, I'm sure that you'll be shocked to discover that Xen's support for hardware virtual machines leverages portions of QEMU's console, serial port, and display technologies. Shocked, I say! I don't know where virtualization would be without Fabrice Bellard (Mr. QEMU), but things wouldn't be as good as they are now.

Creating a hardware virtual machine (HVM) is simpler than creating a paravirtualized virtual machine because you can actually use the installation media for the operating system that you want to run in the virtual machine. In order to do this, your configuration file must identify the virtual disk to which you are installing that operating system, which can be a file, disk partition, logical volume, or any other pre-allocated unit of storage. The virtual disk that you specify in your HVM configuration file will be treated as a disk, not as a disk partition, regardless of where it is physically located, and thus should not include a partition number in the domainU name portion of that disk entry.

The easiest way to create a configuration file for a hardware virtual machine is to simply copy an existing one. The following is a sample HVM configuration file for a 64-bit Microsoft Windows Vista installation on a logical volume:

```
# Vista Ultimate HVM on LV
import os, re
arch = os.uname()[4]
if re.search('64', arch):
    arch_libdir = 'lib64'
else:
    arch_libdir = 'lib'
kernel = "/usr/lib/xen/boot/hvmloader"
builder='hvm'
memory = 384
maxmem = 1024
shadow_memory = 8
name = "VISTA-LV-HVM"
acpi=1
apic=1
vif = [ 'type=ioemu' ]
disk = [ 'phy:/dev/VM/Vista64,xvda,w',
         'phy:/dev/hdb,hdc:cdrom,r' ]
device_model = '/usr/' + arch_libdir + '/xen/bin/qemu-dm'
boot="cd"
sdl=1
serial='pty'
usb=1
usbdevice='tablet'
```

Figure B7-37 shows a Microsoft Windows Vista Ultimate HVM running in an SDL window.

FIGURE B7-37

A legacy operating system running in a Xen HVM



I tend to install HVMs using an SDL console, and then switch them to use a VNC console after the fact. To do this, replace the `sdl=1` entry in the previous file with the following:

```
sdl=0
vnc=1
vncpasswd = 'thisismypassword'
```

After restarting the virtual machine with the updated configuration file, you will be able to connect to its console by using any VNC client to connected to the IP address of the virtual machine. You will need to supply the password that you specified in the `vncpasswd` entry in order to actually connect to the console, at which point you will need to log in there.

Getting More Information About Xen

This section has only scratched the surface of the many things that you can do with Xen, the many commands that are supported by the Xen Management command, and so on. If you're seriously interested in virtualization, I think that you'll find the hypervisor model to be superior to other types of virtualization, and Xen is the only commercial-quality open source hypervisor available today. You will find Xen itself to provide a robust and powerful virtualization solution.

My favorite sources for up-to-date information about Xen on the Internet are the following:

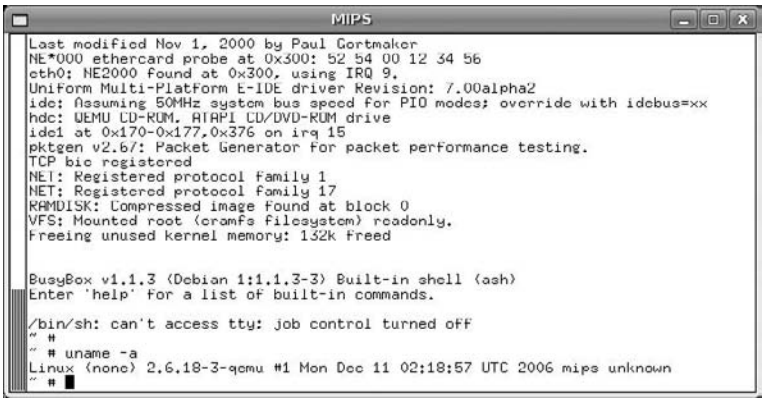
- **Xen on Ubuntu systems:** Ubuntu-oriented information about Xen is available in the Ubuntu Forums at `https://help.ubuntu.com/community/Xen`.
- **Xen Wiki:** An open, editable Wiki found at `http://wiki.xensource.com/xenwiki`. This Wiki contains a tremendous amount of information, some of which is up-to-date. It is the best place to find information about supported/suggested processors that support virtualization, supported/suggested motherboards for use with those processors and Xen, and a huge amount of HOWTO and technical information.
- **Open source Xen mailing lists:** Community mailing lists that you can sign up for at `http://lists.xensource.com`. This site provides access to multiple mailing lists that you can either subscribe to or simply browse or search the list archives. These include lists with focus areas such as standard xen users (xen-users), the Xen developer community (xen-devel), developers of open source tools for Xen (xen-tools), and developers focused on Xen security (xense-devel). Of these, the xen-users mailing list is probably the most useful for most people. As with most mailing lists, you can subscribe to these lists in batch mode, so that you only receive summary digests of recent posts. A real-time subscription to these mailing lists, where you receive every message as it is sent, would probably be quite overwhelming.
- **Citrix customer forums:** Mailing lists specific to the commercial Xen products offered by Citrix, but which are still open to the public. You can register for these at `http://forums.citrix.com` or simply browse their contents, although it would be wrong to pose questions here unless you are a Citrix customer. Again, these forums are specific to Xen from Citrix, but still often contain relevant information for open source Xen users.
- **Xen on other Linux distributions:** Linux distributions such as Fedora and Novell/SUSE, which are both active proponents of Xen, also provide a variety of Xen-related resources. Fedora provides general information about getting started with Xen, Xen tools, and an associated mailing list, which you can find at `https://www.redhat.com/mailman/listinfo/fedora-xen`. Novell provides a "Project Xen Technical Preview" page at `http://forge.novell.com/modules/xfmod/project/?xenpreview`, which contains information about their latest Xen patches and general links to Novell's Xen documentation and other sites.

Installing and Using QEMU

QEMU (`http://bellard.org/qemu`) is an amazingly powerful open source CPU emulator by Fabrice Bellard and a host of others. QEMU can emulate a huge variety of architectures, processors, and related peripheral hardware. QEMU can run operating systems and software that are compiled for platforms including 32- and 64-bit x86, 32- and 64-bit PowerPC, Motorola 68000, 32- and 64-bit SPARC, SH, MIPS, and ARM. QEMU virtualizes a complete hardware environment for each of these processors/architectures, including the CPU and peripheral devices like one or more serial ports, network interfaces, a PCI bus, VGA graphics adaptors, sound hardware, disk drives, USB components, and more. This enables you to run unmodified guest operating systems for any of these architectures/processors in QEMU. For example, Figure B7-38 shows an xterm window in which QEMU has booted a MIPS kernel and is therefore also running operating system binaries that have been compiled for the MIPS platform.

FIGURE B7-38

A MIPS kernel and root filesystem running under QEMU



To install QEMU on an Ubuntu or Kubuntu desktop system, use your favorite package management software (Synaptic on Ubuntu, Adept on Kubuntu, or the command-line tools anywhere), as described in Chapter 23. The core package is simply called `qemu`, which may also install some other applications and libraries on which QEMU depends. You may also want to install the `kqemu-source` and `kqemu-common` packages. The former is the source code for a Linux kernel module that improves the performance of QEMU when emulating an x86 system on an x86 system, and the latter provides the files that are used to load this module when you boot your Ubuntu or Kubuntu system. You will also need to have installed the header files or complete source code for the kernel that your system is running in order to compile the `kqemu` module. The source code for this module is installed in your `/usr/src` directory as a tarball called `kqemu.tar.bz2`. To compile and install this module, do the following:

```
cd /usr/src
sudo tar jxvf kqemu.tar.bz2
cd modules/kqemu
sudo ./configure
sudo make install
```

The `kqemu.ko` module will be installed into the `/lib/modules/KERNEL/misc` directory, where `KERNEL` is the kernel that is currently running on your system. You will need to recompile this module whenever the kernel on your system is updated.

QEMU's hardware emulation is so exceptional that significant parts of QEMU are used to support serial connections, graphics hardware, disk images, and so on in the full virtualization implementations of virtualization packages such as Xen and KVM. In fact, KVM works so closely with QEMU that I'm not going to go into much detail in this section—almost everything in the KVM portions of this chapter also applies to QEMU. For example, you can use QEMU to run virtual machines that were created under KVM—the only difference is that they will not be able to do hardware virtualization under QEMU, and will therefore execute solely through software emulation. Still, not too shabby!

For additional information about QEMU, see:

- **Fabrice Bellard's QEMU page (`http://bellard.org/qemu/`):** The QEMU page hosted by Fabrice Bellard, the primary author of QEMU.
- **QEMU Home page (`http://qemu.org`):** The best starting point for downloads, documentation, or general information about QEMU.
- **QEMU User Documentation (`http://qemu.org/qemu-doc.html`):** The primary documentation for QEMU.
- **QEMU User Documentation and more (`http://fabrice.bellard.free.fr/qemu/user-doc.html`):** Links to the same user documentation as that found at the QEMU home page, but also provides links to technical documentation, internals documentation, porting information, information about the disk image format used by QEMU, and much more.
- **QEMU Developers Mailing List (`qemu-devel@nongnu.org`):** You can subscribe to the mailing list at `http://mail.freesoftware.fsf.org/mailman/listinfo/qemu-devel`.
- **QEMU Users Forum (`http://qemu-forum.ipi.fi/`):** A great collection of forums discussing QEMU, recent patches, the use of QEMU on different platforms, and so on.
- **QEMU IRC Chat (**#qemu channel on Freenode.com**):** A great real-time way of asking questions about QEMU (and hopefully, getting answers as well).
- **QEMU Wiki (`http://kidsquid.com/cgi-bin/moin.cgi`):** An unofficial Wiki that provides a good deal of HOWTO information on using QEMU, QEMU networking setup, and so on.
- **Information on converting disk images to/from QEMU:** `http://en.wikibooks.org/wiki/QEMU/Images`
- **Information on QEMU Networking:** `www.h7.dion.ne.jp/~qemu-win/HowToNetwork-en.html`, `http://en.wikibooks.org/wiki/QEMU/Networking`

Summary

This chapter provided an overview of virtual machine technologies, explaining the differences between them and the advantages (and disadvantages) of each, and highlighting which are available for Ubuntu and Kubuntu systems. It also compared emulation with virtualization and introduced popular emulation packages that are available for Ubuntu and Kubuntu systems. Subsequent sections discussed installing and using VirtualBox, VMware Server, KVM, and Xen, and provided an overview of QEMU.

The next chapter begins the server section of the bonus chapters, discussing how to set up and configure the popular Apache Web server for your Ubuntu or Kubuntu system.

Bonus Chapter 8

Setting Up a Web Server

Most of the significant advances in computing technology have what is known as a *killer app* (killer application)—one significantly unique, powerful, and compelling type of application that draws people to that technology in droves and makes it a part of the computing landscape for the foreseeable future. For personal computers in general, that application was the spreadsheet. For the Apple Macintosh, that application was desktop publishing. And for the Internet, that application was the World Wide Web. Sure, everyone loved e-mail, but the World Wide Web has turned the Internet into a seething pool of e-commerce, personal and technical information, social networking, and who knows what else in the future.

Chapter 11 explained how to use Firefox, the world's best Web browser, to surf the Web. This section explains the flip side of surfing the Web, which is how to set up a Web server so that you can deliver Web pages and other content over the Web to anyone who has access to your server. Most businesses and academic environments today have both externally available and internal-only Web servers. Many people even set up Web servers on their home networks to facilitate Web-based scheduling, document sharing, a central repository for photos, and just about anything else that you can think of.

This chapter begins with a bit of background on the Web, how Web servers work, and a history of Apache, the most popular Web server in existence. Subsequent sections explain how to install Apache, how to configure it, where Web content lives in an Apache server, how to extend and customize the Apache Web server, and how to troubleshoot connectivity or permissions problems.

IN THIS CHAPTER

WWW 101

Introduction to Apache

Installing Apache

Configuring Apache

Troubleshooting your Web server

World Wide Web 101

There are many excellent books about the World Wide Web and HTML, and this book isn't designed to be one of them. However, if you are new to the Web, this section provides some quick history and a sampling of Web buzzwords so that I won't surprise you by using new terms at random.

In 1989, what has become the World Wide Web first entered the world in the mind of Tim Berners-Lee at CERN (Conseil Européenne pour la Recherche Nucléaire), the European Laboratory for Particle Physics near Geneva, Switzerland. The term *World Wide Web* wasn't actually coined until 1990, when Tim Berners-Lee and Robert Cailliau submitted an official project proposal for developing the World Wide Web. They suggested a new way of sharing information between researchers at CERN who used different types of terminals and workstations. The unique aspect of their information sharing model was that the servers would host information and deliver it to clients in a device-independent form, and it would be the responsibility of each client to display (officially known as *render*) that information. Web clients and servers would communicate using a language (protocol) known as *HTTP*, which stands for the HyperText Transfer Protocol.

NOTE **Hypertext is just text with embedded links to other text in it. The most common examples of hypertext outside of the World Wide Web are various types of online help files, where you navigate from one help topic to another by clicking on keywords or other highlighted text. The most basic form of hypertext used on the Web is HTML, the HyperText Markup Language, which is a structured hypertext format that I talk about a little later in this section.**

On the World Wide Web, the servers are Web servers and the clients are typically browsers, such as Firefox, Opera, SeaMonkey, Netscape, Microsoft Internet Explorer, Apple's Safari, and many others, running on your machine. To retrieve a Web page or other Web resource, you enter its address as a Uniform Resource Identifier (URI) in your browser by either typing it in or clicking on a link that contains a reference to that URI. Your browser contacts the appropriate Web server, which uses that URI to locate the resource that you requested and returns that resource as a stream of hypertext information that your browser displays appropriately, and you're off and running!

Today's browsers can understand many protocols beyond HTTP, including FTP (File Transfer Protocol, used to send and receive files), file (used to deliver plain-text files), POP (Post Office Protocol, used to send and receive electronic mail), and NNTP (Network News Transfer Protocol, used to send and receive Usenet News postings). Which protocol you use to retrieve a specific Web resource is encoded into the URI, and is referred to as a scheme in Web nerd terms. A URI specifies three basic things:

```
scheme://host/pathname
```

The *scheme* is one of http, ftp, file, and many more, and specifies how to contact the server running on *host*, which the Web server then uses to determine how to act on your request. The *pathname* is an optional part of the URI that identifies a location used by the server to locate or generate information to return to you.

Web pages consist of a static or dynamically generated text document that can contain text, links to other Web pages or sites, embedded graphics in a variety of formats, references to included documents such as style sheets, and much more. These text documents are created using a structured markup language called HTML, the HyperText Markup Language. (Markup languages were introduced and discussed in Chapter 16.) A structured markup language is a markup language that enforces a certain hierarchy where different elements of the document can appear only in certain contexts. Using a structured markup language can be useful to guarantee that, for example, a heading can never appear in the middle of a paragraph. Like documents in other modern markup languages, HTML documents consist of logical elements that identify the type of each element—it is the browser's responsibility to identify each element and determine how to display (render) it. Using a device-independent markup language simplifies developing tools that render Web pages in different ways, convert the information in Web pages to other structured formats (and vice versa), and so on.

You Say URL, I Say URI . . .

As you may have noticed in the other chapters of this book, different Web-aware applications often use different terms to what you and I might simply think of as "Web addresses." URL (Uniform Resource Locator) is the traditional acronym and term for a Web address, but the acronym and term URI (Uniform Resource Identifier) is actually more technically correct. Another acronym and term that you may come across is URN (Universal Resource Name).

The relationship between these acronyms is the following: a URI is any way to identify a Web resource. A URL is a URI that explicitly provides the location of a resource and the protocol used to retrieve it. A URN is a URI that simply provides the name of a resource, and may or may not tell you how to retrieve it or where it is located.

The bottom line is that most people think of and use the terms URI, URL, and "Web address" interchangeably. If you want to pick one to use, URI is the right term to use.

Introduction to Web Servers and Apache

As mentioned in the previous section, the flip side of a Web browser is the Web server, the application that actually locates and delivers content from a specified URI to the browser. What does a Web server have to do? At the most basic level, it simply has to deliver HTML and other content in response to incoming requests. However, to be useful in a modern Web-oriented environment, a Web server has to do several things. The most important of these are the following:

- Be flexible and configurable to make it easy to add new capabilities, Web sites, and support increasing demand without recompilation and/or reinstallation.
- Support authentication to limit users who can access specific pages and Web sites.
- Support applications that dynamically generate Web pages, such as Perl and PHP, to support a customizable and personal user experience.
- Maintain logs that can track requests for various pages so that you can both identify problems and figure out the popularity of various pages and Web sites.
- Support encrypted communications between the browser and server, to guarantee and validate the security of those communications.

The order of importance of these various requirements depends on whether you are a systems administrator or e-commerce merchant, but all modern Web servers must provide at least these capabilities.

Many different Web servers are available today, depending on your hardware platform, the software requirements of third-party software that a Web site depends on, your fealty to a particular operating system vendor, and whether or not you are willing to run open source software, get additional power, and save money.

As you might expect, the first Web server in the world went online at CERN, along with the first Web browser. These were written and ran on NeXT workstations, not exactly the world's most popular platform (sadly enough). The first test of a Web server outside of Europe was made using a server running at the Stanford Linear Accelerator Center (SLAC) in the United States.

The development focus of Web servers that ran on more popular machines was initially the NCSA (National Center for Supercomputing Applications) Web server, known as *NCSA httpd* (HTTP Daemon). Their development of a freely available Web server paralleled their development of the NCSA browser, known as Mosaic. (See Chapter 11 for information about Mosaic and its place in Web browser history.) When one of the primary developers of NCSA httpd (Rob McCool) left the NCSA, a group of NCSA httpd fans, maintainers, and developers formed to maintain and support a set of patches for NCSA httpd. This patched server eventually came to be known as the *Apache* Web server. Although the official Apache Web site used to claim that the name "Apache" was chosen because of their respect for the endurance and fighting skills of the Apache Indians, most people (myself included) think that this was a joke, and that the name was chosen because the Web server initially consisted of many patches—in other words, it was "a patchy Web server."

Although some Linux distributions still provide Apache 1.*x* servers as an installation option, the Ubuntu repositories only provide the apache2 Web server because its code base and configuration mechanisms are much improved over earlier versions of Apache. Apache 2.*x* is a far superior Web server and where future Apache extension development is going to take place. Any references to Apache in this chapter should be taken to refer to the Apache 2.*x* Web server.

Today, Apache Web servers installed at sites across the Internet deliver more Web content than any other Web server. I forget the name of the second most popular Web server, but it only runs on a single operating system (which is not Linux) and therefore loses conceptually as well as numerically.

Installing Apache

The latest version of the Apache 2.*x* Web server for Ubuntu and Kubuntu systems is provided in the apache2 package in the Ubuntu repositories. You can use any of the package management tools discussed in Chapter 23 to install that package. When installing Apache from the command line using aptitude, I suggest that you use the `-r` option so that recommended packages are also installed. Regardless of whether you're using a command-line or graphical tool, I suggest that you install both the apache2 and apache2-doc packages.

NOTE **If you installed your system from an Ubuntu server CD and chose the Install a LAMP server option, the Apache 2 Web server was installed as part of your LAMP (Linux, Apache, MySQL, and Perl) server installation. You can ignore this installation section and move on to the rest of the chapter (although you may want to install the Apache2 documentation).**

Apache 2 File Locations

This section provides a quick overview of the default locations of the configuration files, binaries, and content associated with the Apache 2 Web server on your Ubuntu system:

- `/etc/apache2`: A directory containing the configuration files for the Apache 2 Web server. The primary configuration file in this directory is the file `apache2.conf`.
- `/etc/apache2/conf.d`: A directory containing local configuration directives for Apache 2, such as those associated with third-party or locally installed packages.

- `/etc/apache2/envvars`: A file containing environment variables that you want to set in the environment used by the `apache2ctl` script to manage an Apache 2 Web server.
- `/etc/apache2/mods-available`: A directory containing available Apache 2 modules and their configuration files.
- `/etc/apache2/mods-enabled`: A directory containing symbolic links to actively enable Apache 2 modules and their configuration files, located in the `/etc/apache2/mods-available` directory. This is analogous to the use of symbolic links to start various processes from the scripts in `/etc/init.d` at different run levels.
- `/etc/apache2/sites-available`: A directory containing files that define the Web sites supported by this server.
- `/etc/default/apache2`: A configuration file that determines whether the Apache 2 should automatically start at boot time.
- `/etc/init.d/apache2`: A shell script that uses the `apache2ctl` utility to start and stop an Apache 2 Web server.
- `/etc/mime.types`: The default MIME (Multipurpose Internet Mail Extensions) file types and the extensions that they are associated with.
- `/usr/lib/cgi-bin`: The location in which any CGI-BIN (Common Gateway Interface scripts) for a default Apache 2 Web server will be installed.
- `/usr/sbin/apache2`: The actual executable for the Apache 2 Web server.
- `/usr/sbin/apache2ctl`: An administrative shell script that simplifies starting, stopping, restarting, and monitoring the status of a running Apache 2 Web server.
- `/usr/share/apache2-doc`: A directory that contains the actual Apache 2 manual (in the `manual` subdirectory). This directory is present only if you've installed the apache2-doc package (as suggested earlier).
- `/usr/share/apache2/error`: A directory containing the default error responses delivered.
- `/usr/share/apache2/icons`: A directory containing the default set of icons used by an Apache 2 Web server. This directory is mapped to the directory `/icons` in your Apache server's primary configuration file.
- `/var/log/apache2/access.log`: The default access log file for an Apache 2 Web server. This log file tracks any attempts to access this Web server, the hosts that they came from, and so on.
- `/var/log/apache2/error.log`: The default error log file for an Apache 2 Web server. This log file tracks internal Web server problems, attempts to retrieve nonexistent files, and so on.
- `/var/run/apache2/apache2.pid`: A text file used by Apache 2 to record its process ID when it starts. This file is used when terminating or restarting the Apache 2 server using the `/etc/init.d/apache2` script.
- `/var/www/apache2-default`: A directory containing the default home page for this Web server. Note that the default Apache 2 Web server does not display the content of this directory correctly—I'll use that as an example of configuring a Web site in the next section.

Some of these directories, most specifically the `/etc/apache2` configuration directory, contain other files that are included or referenced by other files in that same directory.

Configuring Apache

As mentioned in the previous section, the configuration files for the Apache 2 Web server are located in the directory `/etc/apache2`. Configuration files for Web sites that are available in an Apache 2 Web server are located in the directory `/etc/apache2/sites-available`. To actually support a site from your Web server, you must create a configuration file for that Web server in `/etc/apache2/site-available`, and then create symbolic links to that configuration file in the `/etc/apache2/sites-available` directory.

The only Web site that is provided out of the box with a standard Apache 2 installation is its default Web site, which displays the Web page shown in Figure B8-1.

FIGURE B8-1

The default page for a newly installed Apache 2 Web server



NOTE **If you are creating a new Web site and want it to be your Web server's default page, you can simply put your content in the `/var/www` directory, where things would work fine immediately.**

The following is a listing of the file `/etc/apache2/sites-available/default`, which activates the site on this server. (I've added line numbers to make it easier to refer to different entries—they do not actually appear in the file!)

```
1. <VirtualHost *:80>
2. ServerAdmin webmaster@localhost
3.
4.   DocumentRoot /var/www
5.   <Directory />
6.       Options FollowSymLinks
7.       AllowOverride None
8.   </Directory>
```

```
9.   <Directory /var/www/>
10.      Options Indexes FollowSymLinks MultiViews
11.      AllowOverride None
12.      Order allow,deny
13.      allow from all
14.   </Directory>
15.
16.   ScriptAlias /cgi-bin/ /usr/lib/cgi-bin/
17.   <Directory "/usr/lib/cgi-bin">
18.      AllowOverride None
19.      Options +ExecCGI -MultiViews +SymLinksIfOwnerMatch
20.      Order allow,deny
21.      Allow from all
22.   </Directory>
23.
24.  ErrorLog /var/log/apache2/error.log
25.
26.  # Possible values include: debug, info, notice, warn, error, crit,
27.  # alert, emerg.
28.  LogLevel warn
29.
30.  CustomLog /var/log/apache2/access.log combined
31.
32.  Alias /doc/ "/usr/share/doc/"
33.  <Directory "/usr/share/doc/">
34.      Options Indexes MultiViews FollowSymLinks
35.      AllowOverride None
36.      Order deny,allow
37.      Deny from all
38.      Allow from 127.0.0.0/255.0.0.0 ::1/128
39. </Directory>
40.
41. </VirtualHost>
```

The first thing that I want to change here is line 2, which sends any mail directed to the Webmaster for this site to `webmaster@localhost`, which probably doesn't exist on your machine. You can either set up a local alias for Webmaster in your mail server configuration (see Bonus Chapter 9, "Setting Up a Mail Server," for more information) or simply change this to an explicit site-wide address that you've already assigned somewhere. I would change this to `webmaster@vonhagen.org`.

Suppose next that you also wanted to make the Apache documentation explicitly available through the URI `/manual/`. Although there is no such directory, it's easy enough to make an entry in the server's configuration file defining a redirect to an existing directory. So let's make one. Inside the VirtualHost block, create something like the entry for the `/doc/` directory that's shown in lines 31 through 38, but simplify it a bit:

```
1. Alias /manual/ "/usr/share/doc/apache2-doc/manual/"
2. <Directory "/usr/share/doc/apache2-doc/manual/">
3.  Order deny,allow
4.  Deny from all
5.  Allow from 192.168.6 127.0.0.1
6. </Directory>
```

NOTE **Again, the line numbers are not actually present in the file—I'm just using them in the example to make it easier to refer to specific lines.**

The first line defines an alias called `/manual/` that actually points to the directory `/usr/share/doc/apache2-doc/manual/`, which is where Apache's online manual lives. The rest of the lines define who has access to that directory and under what circumstances. Line 2 defines the beginning of directives related to the directory `/usr/share/doc/apache2-doc/manual/`, and line 6 identifies the end of a block of directives for a specific directory. Lines 3, 4, and 5 specify how authentication works. Line 3 says that any statements denying access to the directory are processed before any that allow access to the directory. Line 4 denies all access to that directory, while line 5 allows access to that directory from any host whose first three octets are 192.168.6 (the subnet on which this Web server is running), and from the loopback address for the host. After adding these changes to the file (they must come before the `</VirtualHost>` directive shown in line 39 of the previous example because they are part of the definition for this host on this Web server) you can restart the Web server using the same command as before:

```
$ sudo /etc/init.d/apache2 restart
```

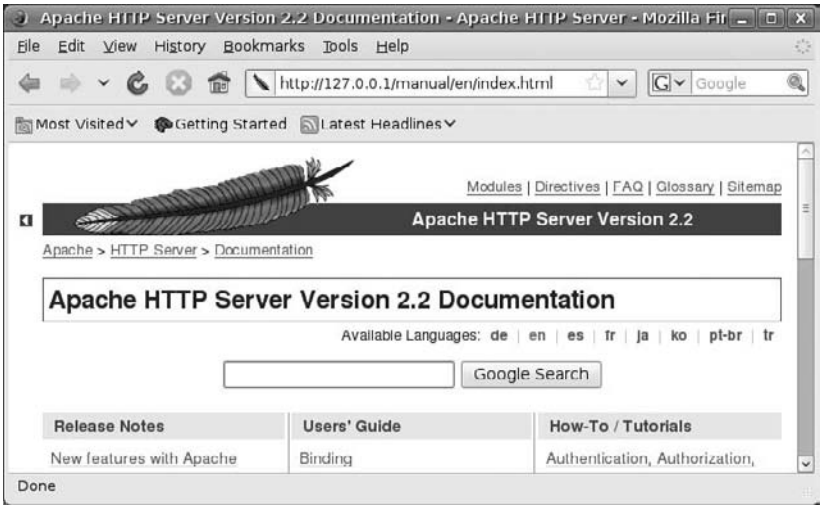
Because the startup script for Apache simply runs the apache2ctl utility (Apache's main administrative utility), you could also restart the Apache server by executing the following command:

```
$ sudo apache2ctl restart
```

Visiting the same URI as before and trying to access the Apache documentation hyperlink now shows the page shown in Figure B8-2.

FIGURE B8-2

The Apache documentation after updating the default server's config file



You may note that there is no equivalent to line 34 of the original server configuration file in the new entry that we created. This is because there was no need to provide these directory browsing options because I knew that the directory contained HTML files, so that the following options were not necessary:

- **Indexes:** Shows an index of the directory if no `index.html` file is present.
- **MultiViews:** Enables content negotiation, where the browser tries to find the best match for a request. In my case, I only want to see the docs in my default language, locale, and character set, so no negotiation is necessary.
- **FollowSymlinks:** Enables the Web server to follow symbolic links. Because I know that there are no symbolic links in this directory, there's no need to specify that they should be followed.

The Apache documentation is really quite good and explains all of the site configuration directives that I don't have space to describe here.

Troubleshooting

As in any debugging or troubleshooting exercise, log files are your friends. Lines 25, 28, and 29 in the original server configuration file shown earlier in this chapter identify the log files used by this server, and the level of logging that occurs. Line 24 identifies the name of the error log file as `/var/log/apache2/error.log`. Line 28 sets the logging level to `warn` (warnings), which is slightly more useful than only logging errors, but is not as useful as `debug` when actually debugging a new site or server. Line 30 tells the server to create a single log file named `/var/log/apache2/access.log` that will log all access requests to the server in NCSA combined log format.

The following Apache log files are exceptionally useful for debugging purposes:

- `access.log`: Shows all attempts to access the server, listing the IP address of the host that attempted access, a timestamp, the actual request that was made, and information about the browser that the request was received from.
- `error.log`: Shows all errors of level warning or above (i.e., more serious) that the server encountered when trying to process an access request. This includes pages that can't be found, directories to which access was denied, and so on.

Apache 2's logging levels are very useful in controlling the amount and type of information that appears in the Apache logs. These levels are as follows:

- **emerg:** Reports only emergency conditions that make the Web server unstable.
- **alert:** Logs situations requiring immediate action, and which may identify problems in the host system.
- **crit:** Logs critical errors that may indicate security, server, or system problems.
- **error:** Reports noncritical errors that indicate missing pages, bad server configuration directives, and general error conditions.
- **warn:** Logs messages that warn of noncritical problems or internal conditions that should be investigated.
- **notice:** Reports normal but significant conditions that should still be looked into.

- **info:** Logs informational messages that may help you identify potential problems or suggest possible reconfigurations.
- **debug:** Logs pretty much every state change on the system, such as every file open, every server activity during initialization and operation, and so on

You should never set the log levels lower than crit on any production Web server. I generally find that warn is the best choice for production servers, and I use notice, info, or debug only when a server is actually having performance or responsiveness problems. Also, don't forget that you actually have to look at the log files for them to be useful.

More Information

You can find many excellent sources of information about Apache 2 on the Internet and in books in which the sole focus is Apache 2. The following are some of my favorite resources:

- `http://httpd.apache.org/docs/2.0/`: It doesn't get much more definitive than the actual Apache 2 docs. However, in addition to reference material, the Apache 2 docs include several tutorials and how-to–style articles that provide practical, hands-on information.
- *Apache Server 2 Bible*, 2nd Edition by Mohammed J. Kabir (Wiley, 2002; ISBN: 0-7645-4821-2).
- *Hardening Apache* by Tony Mobily (Apress, 2004; ISBN: 978-1590593783).

Your favorite search engine, such as Google or Clusty, can help you find information about specific error conditions, warning messages, specific server configurations, and so on.

Summary

This chapter discussed how to install and configure the Apache 2 Web server, one of the most commonly used applications on any networked Linux system. A Web server is not only useful for delivering Web pages, but also provides the foundation of Web-oriented applications that are often used for system administration and monitoring tasks.

The next chapter continues the server-oriented theme of this chapter by discussing the mail servers that are available for Linux systems. The majority of the chapter discusses how to install and configure the Postfix mail server, a powerful, flexible, and easy-to-configure mail server that is a popular alternative to the traditional Linux sendmail server. This chapter also explains how to install and integrate related software, such as anti-spam and virus-scanning software.

Bonus Chapter 9

Setting Up a Mail Server

As discussed in the introduction to Bonus Chapter 8, most of the significant advances in computing technology have what is known as a killer app (killer application)—one significantly unique, powerful, and compelling type of application that draws people to that technology in droves and makes it a part of the computing landscape for the foreseeable future. For the Internet, this application was the World Wide Web, but for networking in general, this application is probably e-mail. I can't deny that the Internet is the world's biggest and most powerful network, but many businesses already depended on e-mail long before they had access to the Internet.

E-mail can be a tremendous time sink, but it is the lifeblood of much business and personal communication nowadays. Some of the many advantages of e-mail for business communication are that it is asynchronous (you can send mail to anyone at your company at any time, and they can read it as soon as they have the chance) and location-independent (you can send mail to anyone at any location where your company has an office). Personal e-mail allows us to stay in touch with friends and family regardless of where they are. However, all of this requires a reliable transport mechanism (the Internet or your corporate network) and reliable software that handles mail transfer between computers, and mail delivery into your mailbox, and provides the ability to read, store, and reply to e-mail.

Chapter 9 discussed the e-mail capabilities of Evolution, a powerful piece of personal information management software that is installed as part of every nonserver Ubuntu installation. This chapter explains how to install and configure Postfix, an equally powerful piece of software that handles mail transfer; how to integrate Postfix with various mail delivery packages; and how to extend its capabilities to filter out the spam that drives many people crazy nowadays. However, before I get into the details, the first section discusses how e-mail works, defines some basic terminology and required software components, provides an overview of available Linux mail servers, and discusses some good reasons for running your own mail server. If you're an experienced sysadmin or are feeling lucky, feel free to skip ahead to the installation, configuration, and antivirus and spam filtering sections.

IN THIS CHAPTER

Reviewing Linux mail servers

Installing Postfix, spam-checking, and virus-scanning packages

Configuring Postfix

Adding POP3 support to your mail server

Adding spam filtering

Adding virus scanning

Introduction to Mail Servers

The components of an e-mail system can be confusing, so I usually generalize them as the difference between receiving and retrieving—my mail server receives e-mail for me, and my e-mail client (Evolution) lets me retrieve my e-mail from there to read it, archive or delete it, and so on. However, in official mailer-speak, there are actually three components involved in a mail system:

- **Mail Transfer Agent (MTA):** A program that handles transferring e-mail from one computer to another, using the Simple Mail Transfer Protocol (SMTP). Well-known examples of MTAs are Microsoft Exchange, Sendmail, Postfix, qmail, and exim. MTAs are often referred to as mail servers, which can be confusing because this is also the term commonly used to identify the host on which an MTA runs.
- **Mail Delivery Agent (MDA):** A program that actually delivers the e-mail received by an MTA into my personal mailbox and can often do both preprocessing, such as spam filtering and virus detection, and auto-responder tasks, such as sending "I'm on vacation" mail. The MDA also often handles the presentation of mail from the mailbox to the user, supporting protocols such as the Post Office Protocol (POP and POP3) and the Internet Message Access Protocol (IMAP) on the host where mail is stored. Common examples of MDAs on Linux systems are procmail, Qpopper, Cyrus, pop3d, Courier IMAP, and `/bin/mail`. MDAs can be directly integrated into an MTA, as is the case with Microsoft Exchange and often with Linux MTAs when mail is read directly from a traditional UNIX-style mail file.
- **Mail User Agent (MUA):** A program that lets you send mail and read the messages in you personal mailbox using protocols such as POP, POP3, and IMAP. Most MUAs also provide support for composing and managing mail messages. Common examples of MUAs are Evolution, Microsoft Outlook, mutt, kmail, pine, and Thunderbird.

In most cases, users never see the MDA because it is integrated with the MTA, either by the software vendor or by the system administrator. Similarly, most users could care less about the MTA and use the term "mail server" to identify the host on which whatever MTA they are using runs. Although the MUA is the most important component of a mail system from the user point of view, the MTA is the most important component from a mail system's and system administrator's point of view because it both performs the actual transfer of mail from one computer to another and identifies the MDA and any other applications involved in preprocessing mail on its way to the user's mailbox.

Popular Linux Mail Transfer Agents

As you might expect, several MTAs are available for Linux systems. Although this chapter focuses on installing and configuring Postfix, other Linux MTAs all have their devotees. For this reason, it's nice to know a bit about each of the most popular ones so that you can feel that you made an informed decision by following the suggestions in this chapter and going with Postfix. The next few sections provide an overview of the most popular Linux MTAs and some of the features, benefits, and downsides of each.

Exim

Exim is a popular, fast, and flexible MTA written by Philip Hazel at the University of Cambridge; it was originally written as a replacement for a popular, older MTA known as smail. Exim is very popular with system administrators who also run mailing lists such as mailman, because it can be configured to automatically look for addresses that are actually mailing lists before delivering mail to end users. It supports an extensive set of features for checking incoming mail for problems without the external applications that most other MTAs require. Exim is the default MTA used on Linux distributions such as Debian.

For more information about Exim, see:

- **Exim Home Page:** `www.exim.org`
- **Exim Documentation and FAQs:** `www.exim.org/docs.html`

Postfix

Postfix is a popular, scalable, and secure MTA written by Witse Venema while working at IBM. Postfix was originally known as VMailer and was even marketed by IBM as Secure Mailer. In 1999, its name became Postfix, and the rest is history. Postfix is the MTA that is installed when you specify that you want to install a mail server when installing an Ubuntu Server distribution.

Postfix is reliable, fast, and easy-to-administer. Its configuration file is easy to read and edit, although you still have to know the many options that you can set and their possible values. It is a drop-in replacement for Sendmail, with the exception that it requires tasks to run as a special Postfix user and group for security reasons.

For more information about Postfix, see:

- **Postfix home page:** `www.postfix.org`
- **Postfix Resources and FAQ:** `www.seaglass.com/postfix`
- **Postfix Shrine:** `www.stahl.bau.tu-bs.de/~hildeb/postfix`

Qmail

Qmail is an extremely fast, secure, and scalable MTA written by Dan Bernstein, a Computer Science professor at the University of Illinois at Chicago who is well known for his interest in cryptography and security. This focus manifests itself in qmail—since 1997, there has been a standing reward of $500 for anyone who finds a verifiable security hole in the latest version of qmail. It has never been claimed.

Qmail is extremely fast, secure, and scalable. Unfortunately, it is also encumbered by strange licensing requirements (for example, you cannot redistribute modified qmail source code without the direct approval of your changes by the author), it is not compatible with the Linux Filesystem Hierarchy Standard (it lives under `/var`), and its initial configuration and the irritations associated with adding new users is responsible for at least some portion of my gray hair. However, the licensing issues have primarily been responsible for its lack of widespread adoption by the open source community, because it isn't open source. As such, several qmail patches are available to add various capabilities, but it is simply tedious to always have to apply a variety of patches.

For more information about qmail, see:

- **Dan Bernstein's qmail page:** `http://cr.yp.to/qmail.html`
- **Life with qmail:** `www.lifewithqmail.org`
- **qmail home page:** `www.qmail.org`

Sendmail

As mentioned earlier in this chapter, Sendmail is the most common mail transfer agent in use on Linux and UNIX systems today. Note that I did not say "the most popular"—Sendmail comes as the default MTA for most UNIX and many Linux distributions, and therefore is often used "because it's

there." If Sendmail is used on many existing systems in your organization, it's also difficult to argue for using a different MTA on any new box, based on the "don't fix it if it's not broken" approach to system administration.

Sendmail was written by Eric Allman, whose delivermail program was the original ARPANET mail delivery system provided with 4.0 BSD UNIX and early versions of 4.1 BSD. However, as the ARPANET grew, and then eventually became the Internet, delivermail proved to be too inflexible, largely because it used compiled-in configuration information. Sendmail was developed to use an external configuration file, and was first delivered with later versions of BSD 4.1. Sendmail's configuration syntax is both tremendously rich and fiendishly arcane, and is the biggest argument against Sendmail. The configuration files actually used by Sendmail are generated from the configuration file maintained by a system administrator using the m4 macro processor, which makes them tremendously painful to debug and maintain. Sendmail configuration questions are tremendously popular during the interview process for system administrators.

On the positive side, no Internet mailer is more mature than Sendmail because it has been in active use for over 20 years. It scales well, and can be integrated with most of the popular anti-spam and antivirus packages that are available. Sendmail has a somewhat undeserved reputation for security problems, which IMHO largely is because it is the biggest and oldest target.

For more information about Sendmail, see:

- **Open Source Sendmail Home Page:** `www.sendmail.org`
- **Sendmail, Inc.:** `www.sendmail.com` (where Eric Allman is CTO)

Why Run Your Own Mail Server?

There are a number of good reasons to run your own mail server, but doing so also introduces both additional responsibilities and internal costs. All domain hosting companies and ISPs will manage a mail server for a domain, giving you some amount of control over its configuration. For example, you are usually given access to a Web-based control panel that enables you to create some number of e-mail accounts and mail aliases. However, this does not give you any control over the internals of the mail server. The following are some of the best reasons for wanting to run and manage your own mail server:

- No limitation on the number of users or aliases. Most ISPs and domain hosting companies provide a certain number of each, and then impose additional charges for higher numbers.
- No limitation on the number of domains that can be served by a single mail server. Most ISPs and hosting companies charge for each domain and require that you administer each domain separately. You can configure a single mailer to support all of your domains, sharing a central configuration across multiple domains.
- The ability to use a different mail server than the one that is provided by your ISP or domain hosting company. You can select the MTA and MDA combination that best suits your requirements, not the convenience of your ISP or domain hosting company.
- The ability to integrate centralized spam and antivirus checking that is under your control and which you can fine-tune at no cost.
- Easy integration of additional capabilities such as mailing lists that may not be supported by your ISP or domain hosting company.

Although deciding to run your own mail server provides several benefits, it introduces several costs and responsibilities that you must be prepared for. Some of the more important of these are the following:

- Security becomes your concern and responsibility. When an ISP or domain hosting firm manages your mail server, its security is their responsibility. When you manage your own mail server, you are directly responsible for its security, which is especially important if you run an IMAP mail server, where all of your personal, SMB, or enterprise mail is permanently stored on the mail server.
- You will need to make sure that any host on the Internet can always contact your mail server. Beyond being responsible for the uptime of the machine that hosts the mail server, you will need to make sure that the MX records for all domains that it supports correctly identify the mail server. It will either have to have a fixed IP address it, or will need to be correctly identified by a dynamic DNS system that is always available.
- You will become responsible for all of your mail server's system administration. In addition to ensuring that it is configured correctly (no open relays, please), you will be responsible for ensuring that a mail server is always available at the location identified by your domain's MX record, that backups are done, that the system and all of the software that you depend on are always kept up-to-date, and that spam filtering and antivirus checking are always up-to-date and correctly configured, and much more.

It's easy to overlook incidental costs when setting up your own mail server. When everything works correctly, as it does most of the time, things are great. But when something breaks, you are the one in the hot seat.

Installing Postfix and Friends

This section explains how to install the packages necessary for a simple Postfix mail server that automatically scans messages for spam and viruses and which users can access via the POP3 protocol. Postfix and related packages are installed in different ways depending on whether you are running a system installed from an Ubuntu Server CD or an Ubuntu or Kubuntu Desktop CD. If you are installing an Ubuntu server distribution, installing a mail server is a suggested option during the installation process, but this installs only the core Postfix package, not all of the other packages that I suggest that you install. (If you do not install a mail server when initially installing an Ubuntu server distribution, you can always subsequently install Postfix and related packages using `apt-get` or `aptitude`.) Otherwise, you can use any of the tools discussed in Chapter 23 (`apt-get`, `aptitude`, the Synaptic Package Manager, or the Adept Package Manager) to manually install Postfix and related packages. Ubuntu and Kubuntu enable you to do things your way.

Packages to Install

In alphabetical order, the packages that I suggest you install are the following:

- `clamav`: An antivirus package that scans incoming mail for attachments that contain a virus and quarantines those messages so that people don't accidentally invoke the problematic attachment.
- `mailx`: A package that contains a simple command-line tool for sending mail, which is used by many system utilities, but (in our case) is extremely useful for quick mail server tests.
- `mailscanner`: A package that scans incoming mail using a variety of spam detection techniques and discards any matching mail.

- `postfix`: The Postfix mailer and all related utilities and configuration files.
- `postfix-doc`: Documentation for the Postfix mailer, installed in `/usr/share/doc/postfix` and `/usr/share/doc/postfix-doc`.
- `postgrey`: A greylisting implementation for use with Postfix. Greylisting causes incoming mail requests to be rejected the first time that they arrive, but accepted when the remote mail system tries to deliver them a second time. Many spammers try to deliver mail only once, so this simple utility can significantly reduce the amount of spam that your site receives.
- `qpopper`: A package that enables users to read their mail remotely, using the POP/POP3 protocol that is supported by most MUAs.

You don't have to install all of these now, of course, but I'm going to assume that these are all installed in subsequent sections of this chapter. I know what they say about what "assume" means, but in this case that won't include me.

Although the way in which you install applications is up to you, I suggest that you use Aptitude to take advantage of its ability to install recommended packages as well as the basic packages required to run and support the Postfix mail server and related applications on your Ubuntu or Kubuntu system. Synaptic and Aptitude will install all required packages, but will not get you all of the bells and whistles associated with some of the recommended packages for Postfix. To install the Postfix mail server and related applications from the command line using Aptitude, execute the following command:

```
$ sudo aptitude -y -r install postfix postfix-doc postgrey \
      mailscanner mailx qpopper clamav
```

Installing all of these packages, their dependencies, and recommended packages, will require around 65MB of disk space. After being prompted for your password, installation will begin without confirmation (because of the `-y` option), and away we go!

Built-in Postfix Configuration Models

Regardless of whether you are installing Postfix and friends from the command line or graphically, you will be required to select a default configuration model during the installation process. Five different configuration models are built into the Postfix package, each of which sets up the Postfix configuration files in a certain basic way. The following are the configuration models:

- **Internet site:** Mail will be sent and received directly from this machine via SMTP using Postfix. This is the default choice, and is probably the right choice if you are configuring an incoming and outgoing mail server for your domain and you have a direct connection to the Internet.
- **Internet site using smarthost:** This machine either receives mail directly from the Internet or from another machine using a utility such as `fetchmail`, but sends mail through another mail server. This is probably the right choice if you want to run a home, SOHO, or SMB network with its own mail server but want to leverage your ISP or domain hosting company's existing mail server.
- **Local delivery only:** The machine only delivers mail locally, and does not send mail over any network. This is the right choice if you are setting up a system on which you want system status messages to be delivered by local e-mail, or a system shared by multiple users who need to exchange mail only between themselves.

- **No configuration:** The installation process will not modify the generic Postfix configuration files. You will have to make extensive manual modifications to them to get Postfix to work correctly on your system.
- **Satellite system:** This machine does not receive mail and forwards all outgoing mail to another mail server for actual delivery.

If you are setting up a system that doesn't match any of the presupplied configurations, it's probably best to select the default value, Internet Site, because this will provide a reasonable template that you can then customize to match your exact requirements.

During the installation process, you will see a screen like the one shown in Figure B9-1 if you are installing Postfix and friends via `aptitude`. (You'll see a similar screen if you're using Synaptic or Adept, but—of course—it will be graphical.)

FIGURE B9-1

Available Postfix configuration models

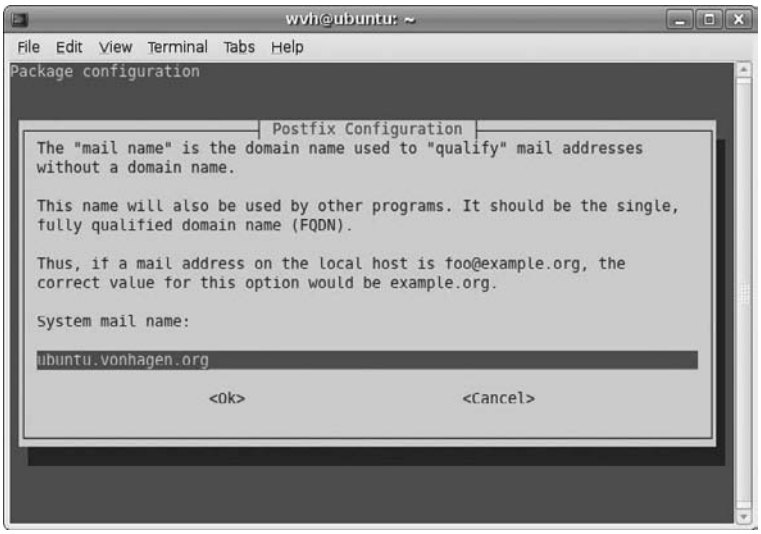


This screen displays the different configuration models that are built into the Postfix package. In most cases, you will want to select the default configuration model, Internet Site, by pressing Tab to move the cursor to the OK option and then pressing Return. You can also use the up/down arrows to select another configuration model, press Tab to move the cursor to the OK option, and press Return to accept the highlighted choice. The screen shown in Figure B9-2 displays.

NOTE **If you simply press Return without tabbing to the OK option, you'll see another screen that provides some additional information about these options. If you continue on from that screen, the screen shown in Figure B9-1 redisplays.**

FIGURE B9-2

Specifying the name of your mail host in text mode



This screen (or its graphical equivalent) enables you to specify the name of this host as any external hosts will look for it, which should therefore be a fully qualified host name that includes the name of your domain. By default, the actual name of the system is displayed. Use the left/right arrow and delete keys if you need to make any changes to the name of this host. Press Return to accept the name on this screen. The remainder of the configuration process will continue with no further questions.

Configuring Postfix

As installed by default, your Postfix mail server is almost ready for prime time. It will already work fine locally, which you can test by using the command-line mail utility (which is provided by the `mailx` package that you installed on the previous section). To make your mail server able to accept incoming requests from and send mail to other hosts on a network, you will need to extend its capabilities somewhat. Similarly, the Postfix server by itself does local mail delivery, which is fine if all of your users want to log in on your mail server and read their mail there. Unfortunately, that is both poor security practice and a hassle for everyone. So you'll need to do a bit of fine-tuning to make our new Postfix mail server more usable for and by everyone.

The next section explores the default Postfix configuration file that was installed and configured in the previous section. The two sections following that address how to make your Postfix server work on your network and how to configure your mail server so that it accepts incoming POP mail requests.

NOTE **As discussed in Chapter 9 and mentioned earlier in this chapter, POP and IMAP are the primary protocols used for remote systems to read e-mail on a mail server. The key difference between them is that POP is designed to store e-mail on the client, while IMAP is designed to store e-mail on the server. Each model has its advantages, but I believe that most people are more familiar with the POP model, so I'm going to focus on that in this chapter.**

Postfix Configuration Files

The Ubuntu Postfix mail server uses three configuration files, all of which are installed in the directory `/etc/postfix`:

- `dynamicmaps.cf`: Identifies additional capabilities to be loaded at run time based on the type of operation that is being performed.
- `main.cf`: Contains configuration information for Postfix that is used during message processing.
- `master.cf`: Contains parameters used when the Postfix master programs runs other programs from `/usr/lib/postfix`.

Postfix can also use standard mailer configuration files such as `/etc/aliases` (which allows e-mail addressed to one user to be automatically delivered to another), but this is not a Postfix-specific configuration file.

The `main.cf` configuration file is the only Postfix configuration file that most users will ever have to modify. As installed on an Ubuntu or Kubuntu system, the `main.cf` configuration file for a system following the Internet site model looks like the following (I've removed some irrelevant comments and white space):

```
#myorigin = /etc/mailname
smtpd_banner = $myhostname ESMTP $mail_name (Ubuntu)
biff = no
append_dot_mydomain = no
#delay_warning_time = 4h
smtpd_tls_cert_file=/etc/ssl/certs/ssl-cert-snakeoil.pem
smtpd_tls_key_file=/etc/ssl/private/ssl-cert-snakeoil.key
smtpd_use_tls=yes
smtpd_tls_session_cache_database = \
    btree:${queue_directory}/smtpd_scache
smtp_tls_session_cache_database = \
    btree:${queue_directory}/smtp_scache
myhostname = ubuntu.vonhagen.org
alias_maps = hash:/etc/aliases
alias_database = hash:/etc/aliases
myorigin = /etc/mailname
mydestination = mail.vonhagen.org, ubuntu.vonhagen.org,
                localhost.vonhagen.org, localhost
relayhost =
mynetworks = 127.0.0.0/8 [::ffff:127:0.0.0]/104
mailbox_size_limit = 0
recipient_delimiter = +
inet_interfaces = all
html_directory = /usr/share/doc/postfix/html
```

In order, these configuration variables do the following:

- `myorigin`: An option for Debian and derivative distributions (such as Ubuntu) that enables you to use the first line of the text file `/etc/mailname` as the name of the mailer. This is actually set later in the default file.
- `smtpd_banner`: Sets the value of the banner that is displayed when an SMTP connection is established. In this case, the banner is constructed by using the value of other variables set in the configuration file.

- `biff`: Determines whether the local e-mail notification service (`biff`) should be used for each user with new mail each time Postfix processes the incoming mail queue. In this case, it is disabled because `biff` notifications can cause a performance drain, and are relevant only on the machine that houses the mail files. Users who use local mail files can enable this for themselves.
- `append_dot_mydomain`: Determines whether Postfix should append a domain name to mail sent to user@host. Nowadays, this is usually handled by the MUA, so this is disabled.
- `delay_warning_time`: Uncommenting this parameter would specify the period of time after which users would be notified if mail that they had sent had not yet been delivered. After all, the remote site could be using Microsoft Exchange, and might be down.
- `smtpd_tls_cert_file`: Identifies the full pathname of the file on this system that holds the certificate used by this machine when sending messages using TLS (Transport Layer Security).
- `smtpd_tls_key_file`: Identifies the full pathname of the file on this system that holds the RSA private key for the Postfix SMTP client when sending messages using TLS (Transport Layer Security).
- `smtpd_use_tls`: Determines whether this Postfix server should use TLS when a remote SMTP server announces STARTTLS support. If the remote server does not announce STARTTLS support, the message is sent in the clear.
- `smtpd_tls_session_cache_database`: Identifies the organization and location of the SMTP server TLS session cache used by the `tlsmgr` daemon.
- `smtp_tls_session_cache_database`: Identifies the organization and location of the SMTP client TLS session cache used by the `tlsmgr` daemon.
- `myhostname`: Identifies the actual Internet hostname of this system. By default, this value is the value returned by `gethostname()`.
- `alias_maps`: Identifies the organization and full pathname of the aliases file used for local mail delivery.
- `alias_database`: Identifies the organization and full pathname of the aliases file used for local mail delivery, and which is updated using the traditional `newaliases` command. This is often the same file as that identified by the `alias_maps` parameter, but need not be.
- `myorigin`: Identifies the name of the host or domain that local mail is assumed to come from and is sent to. On Ubuntu systems, it is the name of a file containing whatever you specified as the mail name when installing Postfix. This is suitable for single-domain installations, but insufficient for Postfix mail servers that support multiple domains.
- `mydestination`: A list of domains for which mail is delivered via local mail delivery. I specified ubuntu.vonhagen.org as the name of my host when I installed Postfix—in most cases, mail.domainname (i.e., mail.vonhagen.org, in my case) would be the standard name to use for a mail server to help keep your sysadmins sane and make it easy to locate/identify your domain's mail server.
- `relayhost`: If outgoing mail must be sent to another mail server for delivery, this parameter identifies that mail server; otherwise, it is empty.
- `mynetworks`: Identifies the networks or specific hosts from which this mail server will send mail. This information is specified in both IPv4 and IPv6 formats. In this case, the mail server will only send mail from the loopback network, which is a problem. This is discussed in the following section, "Identifying Trusted Hosts and Domains."
- `mailbox_size_limit`: Identifies the maximum size of any mailbox on the system. In this case, 0 means that there is no limit.

- `recipient_delimiter`: Identifies the separator used internally by the Postfix server between user names and addresses.
- `inet_interfaces`: Identifies the network interfaces on which the machine can receive mail. In this case, the postfix server will listen on all network interfaces.
- `html_directory`: Identifies the location of HTML files that describe how to install and configure a Postfix server and various Postfix features.

A complete list of Postfix configuration parameters and possible values is available in the files section of the online reference information for the `postconf` command (`man 5 postconf`) or online at locations such as `www.postfix.org/postconf.5.html`.

Identifying Trusted Hosts and Domains

The `main.cf` configuration file in the previous section requires only one change to be a fully functional network mail server —the `mynetworks` parameter must be adjusted to add at least the local network. As installed by default, an Ubuntu Postfix server will listen for incoming requests on all network interfaces, but will only allow hosts on its loopback network to send e-mail through it.

The value of the `mynetworks` parameter can be a single host, an IP address and netmask to indicate a range of hosts or a subnet, or any number of comma-separated hosts or IP address and associated netmasks. The value of the `mynetworks` parameter in the `main.cf` file sets what is known as the *relay policy* for the host. A value must be present, or your mail server becomes what is known as an *open relay*, a mail server through which anyone can send e-mail. Open relays are favorite targets for spammers, because they can abuse them indefinitely.

In this case, my mail server handles mail for all of the systems on my SOHO office, all of which are in the 192.168.6.* address family, so I would change the value of the `mynetworks` parameter to the following:

```
mynetworks = 127.0.0.0/8, 192.168.6.0/8
```

I then either restart the Postfix server (using `sudo /etc/init.d/postfix restart`) or use the `sudo postfix reload` command to cause the active Postfix daemon to pick up my changes to the configuration file. All hosts on my network can now send mail through the Postfix mail server.

Rewriting Addresses in Outgoing Mail

Although all hosts on the local network can now send mail through the mail server, many MUAs send mail using the fully qualified domain name of the host from which the mail was sent. This means that mail from me (wvh) on the host `64bit.vonhagen.org` goes out with an e-mail address of `wvh@64bit.vonhagen.org`. Although I don't particularly care about this, I probably don't have a mailbox there, and therefore people can't reply to such mail.

To correct this problem, you can use the Postfix configuration parameter `masquerade_domains` to identify domains for which Postfix should rewrite outgoing mail addresses so that they simply appear to be from `user@domain`. In this case, mail from `wvh@64bit.vonhagen.org` would go out simply as wvh@vonhagen.org, which is my actual e-mail address. In this case, I would add the `masquerade_domains` setting to my `main.cf` file:

```
masquerade_domains = vonhagen.org
```

I then either restart the Postfix server (using `sudo /etc/init.d/postfix restart`) or use the `sudo postfix reload` command to cause the active Postfix daemon to pick up my changes to the configuration file. All hosts on my network can now send mail through the Postfix mail server without identifying the specific host from which the mail originated.

Accepting Mail for an Entire Domain

As mentioned earlier, the `mydestination` parameter identifies what mail the host will deliver locally. The default Postfix configuration file accepts mail for `localhost.localdomain, localhost`, the actual hostname of the machine, and whatever you specified as the mail name of the machine on the configuration screen shown in Figure B9-2 (if that differs from the actual hostname of your machine). To make your new Postfix mail server a valid final destination for all mail for your domain, you must add the name of your domain to the list of destinations in your `main.cf` file. Because I specified `ubuntu.vonhagen.org` as my mail name, my `mydestination` parameter currently looks like the following:

```
mydestination = ubuntu.vonhagen.org, ubuntu,
                localhost.vonhagen.org, localhost
```

I can either add my domain name to this list or set Postfix's `mydomain` parameter to the name of my domain and then leverage the value of that parameter to provide this information. I would add the mydomain parameter setting and then update the value of the `mydestination` parameter, which together would look like the following:

```
mydomain = vonhagen.org
mydestination = ubuntu.vonhagen.org, ubuntu,
                localhost.vonhagen.org, localhost, $mydomain
```

As with other configuration file changes, I would then either restart the Postfix server (using `sudo /etc/init.d/postfix restart`) or use the `sudo postfix reload` command to cause the active Postfix daemon to pick up my changes to the configuration file.

TIP **When you make the changes suggested in this section, you may also want to change the value of the `myorigin` parameter so that it is set to `$mydomain`. During installation, `myorigin` was set to whatever you specified as the mail name of the machine, which is useful because that value was also used in the values for your `mydestination` parameter, which would correctly handle the case the case where someone sends mail to `root@ubuntu.vonhagen.org`. However, after installation, it is more useful to have `myorigin` be the same as the value of `mydomain`.**

Activating Qpopper for POP/POP3 Support

Qpopper (`www.eudora.com/products/unsupported/qpopper/index.html`) is a popular, free POP/POP3 server from the folks at Qualcomm, who both sell and give away the popular PC and Macintosh MUA, Eudora. I like Qpopper because it supports the vanilla Linux/UNIX mbox format that Postfix uses by default on Ubuntu. (When using the mbox mailbox format, all messages are stored in a single file, and each message begins with the string "From `user@domain`," followed by the timestamp for the message.) I suggested that you install Qpopper in the section on installing Postfix—if you didn't and want to activate POP3 support on your new mail server, please install it now and then proceed with this section.

NOTE **As mentioned before, POP/POP3 and IMAP are alternate protocols for accessing your e-mail, and both can be used with Postfix as long as you install an appropriate server for that protocol on your mail server. Qpopper is fast, lightweight, and well-supported. However, if you want to support both POP/POP3 and IMAP, the Cyrus IMAP (`http://asg.Web.cmu.edu/cyrus/`) and Courier IMAP (`www.courier-mta.org/imap/`) packages are popular and are both available in the Ubuntu repositories. However, configuring the Cyrus IMAP server can be tricky, and the Courier IMAP server supports only the Linux/UNIX maildir mailbox format, where each message is stored in a separate file, which is not the default behavior of Postfix on Ubuntu. Hence Qpopper, which is fast, easy, and "just works."**

Ubuntu's Qpopper package provides a ready-to-run Qpopper daemon that is suitable for use with the `inetd` or `xinetd` Internet services daemons that were discussed in the section of Bonus Chapter 3 entitled "Configuring the svnserve Server." Both of these start various network services in response to specific types of incoming requests. If you have not already installed one of these, I suggest that you install the `xinetd` services daemon. You can then integrate Qpopper with `xinetd` by creating the following file with your favorite text editor and saving it as `/etc/xinetd.d/pop3`:

```
service pop3
{
    port            = 110
    socket_type     = stream
    protocol        = tcp
    wait            = no
    user            = root
    server          = /usr/sbin/in.qpopper
}
```

You can then restart the `xinetd` Internet services daemon using the command sudo `/etc/init.d/xinetd restart`, and Qpopper will be working correctly on your system. You can test that Qpopper is responding correctly to incoming POP/POP3 requests by configuring your MUA to point at your new mail server using the POP/POP3 protocol and checking your mail, or you can do so from the command line by using the `telnet` command to connect to port 100 (the POP/POP3 port) and seeing if you can log in, as in the following example:

```
$ telnet ubuntu.vonhagen.org 110
Trying 192.168.6.182...
Connected to 192.168.6.182.
Escape character is '^]'.
+OK Qpopper (version 4.0.5) at mail starting. <28861.1153954154@ulive>
user wvh
+OK Password required for wvh.
pass **********************
+OK wvh has 1 visible message (0 hidden) in 696 octets.
quit
+OK Pop server at ulive signing off.
Connection closed by foreign host.
```

As you can see, Qpopper responded correctly to my request on the POP/POP3 port, accepted my username and password, and listed the amount of pending mail that I have. Success!

Adding Spam Filtering and Virus Scanning to Postfix

Now that you have a running mail server and are receiving incoming connections from other mail servers, you will quickly find that an incredible amount of the mail that you receive is spam (also known as *UCE*, Unsolicited Commercial E-mail). Some of that mail is going to contain virus-laden attachments. You may also, unfortunately, receive a virus from a friend who forwards you something that they thought was cute, but which actually contains a virus payload. If you're lucky, whoever was hosting your mail before you switched to your own mail server was filtering much of this out, and you may not even have known how bad things are in e-mail-land. Now that you're hosting your own mail server, you will quickly see that spam and virus-laden mail is an epidemic. No wonder it takes so long for me to connect to Amazon.com from the office some days…

Aside from the fact that you're now aware of the spam/virus epidemic, it is also now your responsibility to do something about it because you're responsible for your mail server. Luckily, there are some great ways to configure Postfix to automatically handle much of this for you. If you installed all of the packages that I suggested back in the section "Installing Postfix and Friends," you've already installed the postgrey greylisting service and MailScanner. MailScanner scans incoming mail and rejects the spam that it can identify, and then can invoke other packages, such as SpamAssassin (http://spamassassin.apache.org), another package that does an even better job of spam detection, and ClamAV (www.clamav.net), which is an excellent open source antivirus package. All that remains to do now is to integrate these packages and Postfix.

At this point, you may be wondering why you should bother with an antivirus package on Linux because Linux systems are essentially virus-free. I suggest that you integrate an antivirus package as a matter of course because some of the users of your mail server may also want to read mail on the Microsoft Windows machines that they otherwise use for playing games. More seriously, you may be setting up a mail server in a SOHO, SMB, or enterprise environment in which many of your users need to use Microsoft Windows systems. Either way, installing an antivirus package when you configure your mail server is preventative medicine. It costs you nothing, and if it saves one person the agony of getting a virus on their system and having to try to remove it, you're way ahead.

The next three sections discuss various levels of spam and antivirus protection. The first discusses greylisting as implemented by the `postgrey` package. The second discusses some simple Postfix parameters that you can add to your `main.cf` configuration file that will reject much incoming spam immediately, saving MailScanner the trouble of checking it. The third section then explains how to integrate MailScanner with Postfix and ClamAV with MailScanner to identify as much spam as possible that comes from correctly configured mail servers, and to check for viruses at the same time.

Greylisting via postgrey

In the section on "Installing Postfix and Friends," I suggested that you install `postgrey`, which is a greylisting application that was designed to work with Postfix. That package is set up and started automatically and requires no other mandatory configuration—if you installed it when you installed and started Postfix, it's already working for you.

As mentioned in the installation section, greylisting is a cool technique wherein your mail server rejects all incoming mail the first time that a remote mail server tries to deliver it to you. Most spammers try only once and then move on—they can't afford to endlessly retry sending mail to hosts that have already rejected it once. All real mail servers that are sending real mail will try to deliver the mail again, at which point your Postfix server will actually talk to them. At that point, the other techniques described in this section kick in.

NOTE **If you're interested, the term *greylisting* is related to two other techniques for avoiding spam. The term *blacklisting* refers to the practice of refusing to accept mail from hosts that are on a list of known or suspicious hosts. The term *whitelisting* refers to the practice of accepting mail only from a short list of known hosts, specific people, and so on. It's a shame that we all have to resort to these sorts of techniques to protect ourselves from idiot spammers, but that's the reality of the Internet today.**

Although `postgrey` requires no specific manual configuration, there are several command-line options that are not active by default but which you may want to take advantage of. The `postgrey` service also has several features that you may want to take advantage of to customize how greylisting is done on your system. You can whitelist specific sites by adding them to the various whitelist files in `/etc/postgrey`. For example, the file `/etc/postgrey/whitelist_clients` identifies sites for which no greylisting will be done—all incoming mail from these sites will be accepted immediately. Similarly, the file `/etc/postgrey/whitelist_recipients` identifies users of your mail server that do not want greylisting—in other words, users for whom incoming mail will be accepted immediately, regardless of the site from which it originates. You should read the online man page for postgrey using the `man postgrey` command to see if there are options that you want to take advantage of to fine-tune how the `postgrey` service works on your system.

Adding Postfix Parameters to Reject Bogus Mail

As mentioned in the section "Postfix Configuration Files," a tremendous number of Postfix configuration parameters are available, relatively few of which are actually set in the default `main.cf` configuration file installed with the Ubuntu Postfix package. In general, that's a good sign, letting you know that the defaults are typically correct. However, several of these parameters provide additional capabilities that can help your Postfix mail server reject spam out of hand by checking for things that do not conform to Internet electronic mail standards (as defined in the RFCs). Although the next section, "Integrating MailScanner, SpamAssassin, and ClamAV with Postfix," explains how to scan the content of incoming e-mail for spam, scanning incoming mail takes some amount of time, which reduces the overall performance of your system slightly. Any mail that you can reject before you even have to scan it is therefore a win for everyone.

Adding the following block of commands to your Postfix `main.cf` configuration file can be very effective in reducing spam. (Note that the line numbers are present to make it easy for me to refer to different lines—they should not be present in your `main.cf` file!)

```
1  smtpd_delay_reject = yes
2  smtpd_helo_required = yes
3  smtpd_helo_restrictions =
4     permit_mynetworks,
5     reject_invalid_hostname,
6     reject_non_fqdn_hostname,
7     permit
8  smtpd_data_restrictions =
9     reject_unauth_pipelining,
10    permit
11  smtpd_sender_restrictions =
12    permit_mynetworks,
13    reject_non_fqdn_sender,
14    reject_unknown_sender_domain,
15    permit
16 smtpd_recipient_restrictions =
17    permit_mynetworks,
```

```
18    reject_unknown_recipient_domain,
19    reject_unauth_destination,
20    permit
```

Line 1 is required to use any of the subsequent blocks of restrictions, telling your Postfix server to wait until it receives certain commands from the remote mail server (i.e., a remote SMTP daemon that is a client of your mail server) before rejecting mail based on the rule chains that follow. The `smtpd_delay_reject` phrase is required because some SMTP servers misbehave if you reject mail before they've had the chance to waste at least some of your time. Line 2 rejects mail from any mail server that fails to correctly identify itself using the HELO or EHLO phrases required by the SMTP RFC. This saves you from having to check any of the subsequent rules.

Lines 3 through 7, 8 through 10, 11 through 15, and 16 through 20 are chains of rules that are applied to different portions of your mail server's interaction with a remote mail server:

- Lines 3 through 7 identify restrictions on the HELO or EHLO commands which are used by a remote mail server to identify itself to your Postfix mail server. Line 4 immediately permits connections from any mail server listed in the Postfix `mynetworks` parameter. Line 5 rejects connections from remote mail servers that identify themselves using a malformed hostname. Line 6 rejects connections from remote mail servers that do not identify themselves using a fully qualified hostname. Line 7 says that your mail server will accept HELO or EHLO commands from mail servers that have not failed the previous tests, and allows you to move on to the next set of tests.
- Lines 8 through 10 identify restrictions on the way in which the remote mail servers sends commands to your mail server. Line 9 tells your server to reject connections that send commands to your mail server too quickly, or at least before your mailer has told the remote mail servers that it can accept a fast stream of SMTP messages. Many bulk e-mailers automatically do this to try to speed up deliveries—after all, they have a zillion targets to send mail to, so they can't afford to waste time (except for yours). Line 10 says that your mail server will accept additional data from mail servers that have not failed the previous tests, and allows you to move on to the next set of tests.
- Lines 11 through 15 identify restrictions on the sender of the e-mail that your server is receiving (identified using the SMTP MAIL FROM command). Line 12 immediately permits connections from any mail server listed in the Postfix `mynetworks` parameter. Line 13 rejects mail from any sender whose name is not identified using a fully qualified hostname. (Note that mail from your local network may use a short hostname, but we've already permitted mail from them.) Line 14 rejects mail from senders with unknown domains. Line 15 says that your mail server will accept mail from senders that have not failed the previous tests, and allows you to move on to the next set of tests.
- Lines 16 through 20 identify restrictions on the recipient of the e-mail your server is receiving (identified using the SMTP RCPT TO command). Line 17 immediately permits connections from any mail server listed in the Postfix `mynetworks` parameter. Line 18 rejects mail where the current mail server is not the final target of the mail and that target is not a valid domain. Line 19 rejects incoming mail if the recipient is not in the domains serviced by the current mail server (identified in the `mynetworks` parameter) or is not in the domains that your mail server will relay mail to (identified in the `relayhost` parameter). This prevents your mail server from being used as an open relay. Line 20 says that your mail server will accept mail for recipients that have not failed the previous tests, and allows your mail server to move on to actually processing and delivering the incoming mail.

Postfix provides many other checks that I haven't used here, such as checks that test whether the incoming mail is from a mail server that is known to be a source of spam (i.e., is on a blacklist). In my experience, this type of lookup is slow and these lists are often wrong. For example, some sites that support these blacklists identify any IP addresses that are dynamically assigned by providers such as Comcast as sources of spam. Although this may sometimes be the case, it often is not—for example, Comcast is my ISP, and I promise that I don't send spam.

The sample commands shown in this section should significantly reduce the amount of incoming spam that you receive. Once you've had the chance to experiment with this, you may want to try adding other, similar commands. For a complete list of other checks that you can do to reject incoming spam, see the files section of the online reference information for `postconf` command (`man 5 postconf`) for more information, which is also available online at locations such as `www.postfix.org/postconf.5.html`.

Integrating MailScanner, SpamAssassin, and ClamAV with Postfix

MailScanner is a fast, powerful tool that scans incoming mail for spam and can optionally be configured to scan incoming mail for viruses. The Ubuntu MailScanner package requires that the SpamAssassin package (discussed in the section of Chapter 9 entitled "Automatically Checking for Junk Mail") be installed on your system so that it can invoke the SpamAssassin package after performing its own spam checks, to maximize the chances that it will catch and flag incoming spam.

NOTE **The MailScanner package always claims that it is scanning for viruses even if it is not. This is because its configuration file contains a setting for the virus scanner that you want to use, which by default is set to none. Don't worry; I'll explain how to fix that.**

I suggested that you install MailScanner and the ClamAV virus-scanning package in the section of this chapter entitled "Installing Postfix and Friends." If you didn't, please go back and see that section for information on how to install those packages.

Once installed, it is remarkably easy to add MailScanner to the Postfix food chain:

1. Add the following line to the end of your Postfix `main.cf` configuration file (`/etc/postfix/main.cf`):

   ```
   header_checks = regexp:/etc/postfix/header_checks
   ```

2. Use your favorite text editor to create the file `/etc/postfix/header_checks` with the following contents:

   ```
   /^Received:/ HOLD
   ```

 This line tells Postfix to scan every incoming message for instances of the string "Received:" at the beginning of a line and to move any matching messages (which will be all of them) to a HOLD queue. The HOLD queue is a temporary staging area that you'll use for scanning the incoming mail. Save this file and proceed to the next step.

3. Using your favorite text editor, bring up the MailScanner configuration file (`/etc/MailScanner/MailScanner.conf`) so that you can configure MailScanner to work correctly with Postfix as installed on your Ubuntu system. (By default, the Ubuntu MailScanner package is configured as it would be on a Debian system and is therefore configured for use with the exim mail server.)

4. Find the `Run As User` entries, remove the hash mark at the beginning of the line that says `Run As User = postfix`, and make sure that there is a hash mark at the beginning of any other line that begins with `Run As User`. These entries should now look something like the following:

```
# Run As User = mail
Run As User = postfix
# Run As User = Debian-exim
```

5. Find the `Run As Group` entries, remove the hash mark at the beginning of the line that says `Run As Group = postfix`, and make sure that there is a hash mark at the beginning of any other line that begins with `Run As Group`. These entries should now look something like the following:

```
# Run As Group = mail
Run As Group = postfix
# Run As Group = Debian-exim
```

6. Find the `Incoming Queue Dir` entry and change its value to `/var/spool/postfix/hold`. This entry should now look like the following:

```
Incoming Queue Dir = /var/spool/postfix/hold
```

7. Find the `Outgoing Queue Dir` entry and change its value to `/var/spool/postfix/incoming`. This entry should now look like the following:

```
Outgoing Queue Dir = /var/spool/postfix/incoming
```

8. Find the `MTA` entry and change its value to `postfix`. This entry should now look like the following:

```
MTA = postfix
```

9. Find the `Virus Scanners` entry and set its value to `clamav`. This entry should now look like the following:

```
Virus Scanners = clamav
```

10. Save the `/etc/MailScanner/MailScanner.conf` file and proceed to the next step.

11. Use the following commands to ensure that the Postfix user and group own MailScanner's working directory and the directory where it stores (quarantines) files that are suspected of containing viruses, so that the Postfix mail server can write to these directories:

```
sudo chown -R postfix:postfix /var/spool/MailScanner
```

12. Using your favorite text editor, modify the file `/etc/default/mailscanner` and remove the hash mark at the beginning of the line containing the `run_mailscanner` variable. This line should look like the following:

```
run_mailscanner=1
```

Enabling this variable will cause your system to automatically start MailScanner the next time you boot your system. Save the file and exit from the text editor.

Now, start MailScanner for the first time by using the following command:

```
sudo check_mailscanner
```

Now either restart the Postfix server (using `sudo /etc/init.d/postfix restart`) or use the `sudo postfix reload` command to cause the active Postfix daemon to pick up your changes to the configuration file.

To test that MailScanner is correctly integrated with your mail server, send yourself a piece of mail through your mail server. If MailScanner is working correctly (and the mail was not spam and did not contain a virus, you should see the following banner at the bottom of the message:

```
This message has been scanned for viruses and
dangerous content by MailScanner, and is
believed to be clean.
```

Congratulations! You are now doing a decent amount of spam filtering and are also checking for viruses in any mail received by your mail server!

TIP **As mentioned previously, SpamAssassin is preconfigured for use with MailScanner and comes with several rules that it uses to test whether incoming messages should be marked as spam. You can obtain many other sets of rules for use with SpamAssassin that extend the basic capabilities that it provides. After you're comfortable with the fact that things are actually working correctly, you may want to add some of these other rule sets. Sites such as the Exit0.us SpamAssassin Wiki (`www.exit0.us/`) and the SpamAssassin Rules Emporium forums (`www.rulesemporium.com/forums/`), or a mailing list (`http://lists.maddoc.net/mailman/listinfo/sare-users`) are great sources for additional rules that you may find useful. However, remember that every additional rule that you add increases the amount of time that it takes to check each piece of incoming mail. Don't get carried away, or your mail server may be able to process only one or two incoming messages per day.**

More Information

This chapter provides a thorough explanation of setting up a simple Postfix mail server on your Ubuntu system. If you are going to run your own mail server, you should purchase one of the various Postfix books that are available so that you have an exhaustive, dedicated resource at your fingertips if anything ever goes wrong. The following are some of my favorites:

- ***Postfix* by Richard Blum (Sams, 2001; ISBN: 0672321149):** I have used this book for years, and it has always been an excellent basic reference for me, although it is the oldest and therefore is missing all of the configuration parameters that have been added since 2001.
- ***Postfix: The Definitive Guide* by Kyle D. Dent and Wietse Venema (O'Reilly Media, 2003; ISBN: 0596002122):** Things don't get much more definitive than the O'Reilly book, because Witse Venema is one of its authors.
- ***The Book of Postfix: State-of-the-Art Message Transport* by Ralf Hildebrandt and Patrick Koetter (No Starch Press, 2005; ISBN: 1593270011):** Ralf Hildebrandt is a well-known Postfix expert and advocate—his Postfix Shrine Web site was mentioned as an excellent Postfix resource (more useful if you happen to speak German). Don't worry monolinguals: the book is even better, and it's in English.

As with all hi-tech topics nowadays, your favorite search engine (`www.clusty.com` or `www.google.com` are two popular ones) is your best friend when looking for immediate help with a specific configuration problem, error message, or general Postfix problem.

Summary

E-mail is one of the killer applications of today's Internet, and running your own mail server is standard operating procedure in many enterprise and academic environments, and is getting more and more common for home users. The complete control that running your own mail server provides can be a valuable asset in almost any environment, reducing if not eliminating spam and malicious system attacks through e-mail viruses and Trojans.

This chapter provided background information about the most popular Linux mail systems in use today, and then focused on installing and configuring the Postfix mail server. This chapter also explained how to integrate add-on services such as spam filtering and virus detection into a standard Postfix installation. Postfix is a powerful and easily configured mail server, with which a variety of additional open source services can be easily but tightly integrated.

Bonus Chapter 10 discusses how to install and configure a Dynamic Host Control Protocol Server (DHCP), which enables you to manage and monitor the distribution of IP addresses to client computer systems on your network. Dynamically assigning e-mail addresses enables you to serve a large pool of client systems with a relatively small set of IP addresses, and also simplifies the process to ensure that all client computer systems share the same network configuration data. Using DHCP also simplifies updating basic network information on client systems because this information is stored and updated centrally and can easily be pushed to all clients by the DHCP server.

Bonus Chapter 10

Setting Up a DHCP Server

A long time ago when even I was young, the Internet (or the ARPANET as it was called in those days) was a small place with a limited number of hosts, all of which had fixed IP addresses and names that were maintained in a file that everyone shared with everyone else, and then added their own local modifications for any private, local hosts and networks. However, as the ARPANET grew and more and more hosts became networked, maintaining all of this information in a single file became not only impractical, but also silly. This led to a flexible, software-based service known as DNS (the Domain Name System), which provided a flexible mechanism for identifying the hostname associated with an IP address, and vice versa. (For more information about DNS, see Bonus Chapter 11, "Setting Up a DNS Server.")

Why begin a chapter on DHCP (Dynamic Host Configuration Protocol) by discussing DNS? Because DHCP and its predecessors were largely developed to address similar sorts of scalability and flexibility by providing a similarly flexible service that could dynamically provide various types of network-related information. When few hosts at a site are networked and those hosts are all located on the same subnet, assigning them static IP addresses and statically maintaining other IP-related information such as name-server identities is reasonable. However, as the number of hosts on a network grows and hosts are added more and more frequently, requiring that someone from IT manually install and configure each host's networking quickly becomes impractical. Similarly, networked laptops are frequently used in multiple locations, each of which may require a different network configuration.

This chapter provides a bit of background information on DHCP and its predecessors, the types of problems that they were developed to solve, and then explains how to set up your own DHCP server and some of the issues to consider when doing so. This chapter focuses on DHCP servers for traditional IPv4 networks, but also provides occasional IPv6-related tips.

IN THIS CHAPTER

DHCP overview

Configuring a DHCP server from the command line

Configuring a DHCP server graphically

Troubleshooting DHCP

Overview of DHCP

DHCP stands for Dynamic Host Configuration Protocol. But what does that mean for the users of a network? It means that when everything is working correctly, you should be able to connect your computer to the data jack on the wall with an Ethernet cable, and the University's computer will recognize your computer and allow it to join the network.

DHCP is used to provide network configuration information to the computer systems on a network that are configured to retrieve information dynamically. These systems, known as DHCP clients, locate a system running a software service known as a DHCP server and obtain network configuration information from that system.

DHCP was created by the Dynamic Host Configuration Working Group of the Internet Engineering Task Force (IETF). It was originally created to satisfy the requirements of Internet Request for Comments (RFC) document 1531, which begat RFC 1533 (which begat RFC 2132), which begat RFC 1534, which begat RFC 1541, which begat RFC 2131, and so on. As with all Internet RFCs, reading these documents can provide some fascinating insights into the thought processes that led to the development of many of the concepts and services that we all take for granted on today's networks and the Internet in general. You can access the IETF's RFC archive through the URL `www.ietf.org/rfc.html`, or find those same RFCs at a somewhat easier-to-browse site at `www.faqs.org/rfcs`. Some of these are true classics, such as RFCs 1149, 2549, and so on.

DHCP is based on an earlier protocol for retrieving the information required to boot specific hosts known as BOOTP, and maintains some backward compatibility with BOOTP—hosts that require BOOTP (such as some embedded systems) can retrieve the required information from a correctly configured DHCP server. The key difference between the two is that BOOTP was designed to provide specific responses to requests from specific hosts, whereas DHCP is a much more flexible system that responds to requests from essentially random clients. DHCP also provides a configurable mechanism for reallocating IP addresses to hosts that already have them and recovering network addresses that are no longer in use through its concept of limited-duration leases.

A DHCP client locates and retrieves network configuration information from a DHCP server through the following series of events. (For simplicity's sake, this example ignores routing DHCP requests through relay agents.)

1. A client broadcasts a DHCP discover message to the address `255.255.255.255`. This broadcast includes information that uniquely identifies the client, such as its Media Access Control (MAC) address.
2. Any DHCP servers that receive the offer return an appropriate IP address by broadcasting a DHCP offer message. The IP address that is offered is based on the network interface and subnet from which the request was received, and is either a preallocated IP address that is associated with a specific MAC address, or one from the set of available addresses that this server can offer on the designated subnet. (The set of IP addresses that a specific DHCP server can provide is typically known as the pool of addresses for that server.) IP addresses that cannot be issued by a DHCP server are those that are already in use or those that are reserved for use by clients with specific MAC addresses.

 When broadcasting an offer message, the server temporarily reserves this address so that it does not offer it twice. The information returned to the client also includes information about any other configuration services that the DHCP server can provide for the client.

3. The DHCP client selects the best offer that it has received based on the type and number of services offered and broadcasts a request that identifies the IP address of the server whose offer it wants to accept. This broadcast enables all "also-ran" DHCP servers to drop any addresses that they may have temporarily reserved and move on.
4. The DHCP server identified in the second broadcast message allocates the IP address that it has reserved for the client and enters that IP address into its database of allocated IP addresses, along with a timestamp identifying when it was allocated and an identifier for the time period for which assigned IP addresses are valid (the lease period). The DHCP server then sends an ACK response back to the client that contains all of the network configuration information that it can provide.
5. The DHCP client uses the configuration information sent to it by the DHCP server to configure its network interfaces, DNS servers, and so on, and continues booting. The client also notes the duration of the lease period, and begins a countdown timer that will notify the DHCP server when the lease expires, asking for an extension. When this occurs, and if the extension request adheres to the DHCP server's lease policy, the DHCP server extends the lease.

When a DHCP server issues a lease for an IP address, it either begins a timer or simply polls to verify that the lease period for a specific IP address has not expired. Most DHCP clients notify a DHCP server that they no longer need an IP address as part of their orderly shutdown process. DHCP clients that "terminate abnormally" (i.e., crash) don't have time to do this, but any IP address that they allocated will automatically be freed for reuse by the DHCP server some period of time after the lease expires.

As you can see from this explanation, DHCP is an extremely clever and powerful mechanism that serves both the needs of hosts that require specific configuration settings (and which are therefore identified in the DHCP server's configuration information) and the requirements of random clients that just happen to appear on the subnet that a DHCP server is monitoring.

NOTE **The flip side of flexibility is the ease with which the wrong thing can happen. For example, DHCP servers that are misconfigured or started by accident may provide DHCP clients with completely incorrect information. In my experience, the most common cause of DHCP problems is "randoms" starting invalid DHCP servers on networks that already have an official DHCP server. In enterprise environments, few things can make a sysadmin crankier than someone starting a misconfigured DHCP server on their network(s). This is easy enough to do, especially on multi-homed systems (systems with multiple Ethernet interfaces) where you might have forgotten to identify the Ethernet interface that you wanted your DHCP server to listen on.**

Now that we're all one with the concept of DHCP servers and the types of information that they provide, let's install and configure one.

Installing a DHCP Server

Most server processes aren't installed as part of a default Ubuntu or Kubuntu installation because, let's face it, not everybody needs (or wants) to run servers. However, as with all software packages on Ubuntu and Kubuntu, the package management software discussed in Chapter 23 makes it easy to install a DHCP server. The DHCP server that I suggest installing is the Institutional Service Component's (ISC) v3 server, which is found in the package `dhcp3-server`. This is the latest and greatest DHCP server available, and is certainly the one that provides the most functionality.

If your system has a GUI, I also suggest that you install the `gdhcpd` package on GNOME-based systems, which provides an easy-to-use graphical interface that simplifies configuring and customizing a DHCP server. Unfortunately, I'm not aware of any graphical tools for DHCP configuration on Kubuntu systems—the only graphical KDE-based tool that I have used in the past, `kcmdhcpd`, seems to have vanished long ago. You can always install `gdhcpd` on your Kubuntu system if you are also willing to install all of its GNOME requirements. Installing a graphical DHCP server configuration tool is, of course, optional, but these tools can be quite convenient if you are new to DHCP configuration: They prevent most syntax errors because they create the DHCP configuration file based on fields that you fill in. They can also be useful if you are simply not a fan of using a text editor for everything. Another solution is to install a Web-based system configuration tool such as `webmin` if you are also running a Web server on the system on which your DHCP server is running. For more information about `webmin`, see Bonus Chapter 8.

If your system obtains its address through DHCP, packages such as `dhcp3-client` and `dhcp3-common` will already be installed on your system. The former contains the DHCP client that enables a system to get a DHCP address and related files, and the latter contains documentation, sample files, and utilities that are shared by all of the ISC DHCP packages. Be very careful when installing a new DHCP server in this scenario because this means that another DHCP server must already be running somewhere on your network. Your system will get its IP address earlier in the system startup process than when your DHCP server is started, but this is still a potentially dangerous thing to do.

TIP **If you are also running a DNS server, you may also want to install the `autodns-dhcp` package, which can automatically update a DNS server with information about systems whose IP information was newly assigned via DHCP.**

Alternate DHCP Servers

The Ubuntu repositories provide several DHCP servers in addition to the one that I suggested installing. One or the other of these may be better suited to your network configuration. Also, thanks to the fact that DHCP is indeed a standard, all DHCP servers use the same configuration commands. They may use different configuration files, but the commands that these files contain are the same across all DHCP servers. The following are optional DHCP servers available for Ubuntu and Kubuntu systems are the following:

- `dnsmasq`: A small-footprint DNS server (and proxy) that can also act as a DHCP server. I don't recommend this package as your primary DHCP server unless you have a specific reason not to use BIND, which is the standard, and most popular, DNS server. However, with its DHCP server support disabled, `dnsmasq` can be quite useful on laptops because of its DNS caching capabilities. See Bonus Chapter 11, "Setting Up a DNS Server," for more information about BIND and `dnsmasq`.
- `udhcpd`: A very small-footprint DHCP server designed primarily for use in embedded systems. You may want to install this DHCP server if memory or storage resource minimization is critical.
- `wide-dhcpv6-server`: A DHCP server for IPv6 hosts, this DHCP server requires that you use an IPv6 DHCP client such as wide-dhcpv6-client on any IPv6 systems that you want to configure using DHCP.

The DHCP server that you want to use is up to you. I suggest the dhcpd3-server package because it corresponds to the latest and greatest DHCP server available from the ISC. Your mileage may vary (or your environment may suggest or mandate a different DHCP server).

When running a DHCP server on a system that is itself a client of some other DHCP server, you must make sure that your new DHCP server hands out a different range of IP addresses than the existing DHCP server (and I hope that you're the sysadmin of that network!), is restricted to a different Ethernet interface than the one that your system uses to obtain its IP address, is restricted to specific hardware addresses, or all of the above. It's somewhat unusual to run a DHCP server on a system that itself gets its address via DHCP unless that system has multiple Ethernet interfaces and thus serves as a gateway or router between multiple networks. Most servers have fixed IP addresses so that you can always contact them even if services such as DNS are unavailable. As mentioned previously, setting up a DHCP server that hands out addresses on a network that you are not responsible for is usually bad form, at best.

Managing a DHCP Server from the Command Line

There are always a few people who would rather fight than switch, which in this case means that they would rather create a DHCP server configuration file with a text editor than use a whizzy graphical tool. That's fine—using a text editor is probably faster if you know what you're doing (or are willing to accept the consequences). The next few sections explain how to manually create DHCP server's configuration file and discusses some of the more important options that you may need to add to a DHCP server's configuration file to maximize its usability in your environment.

Creating DHCP Configuration Files Using a Text Editor

This section explains the basics of a DHCP configuration file and the most common set of entries that you will need to put in your configuration file to serve a specific range of IP addresses via a shiny new DHCP server.

TIP **The file `/etc/dhcp3/dhcpd.conf` is the configuration file used by the dhcpd3 server that I suggested installing in the previous section. If you installed another DHCP server, shame on you but I'm sure that you had your reasons, and its configuration file is probably `/etc/dhcpd.conf`, which is the traditional location for a DHCP server's configuration file. An easy way to verify the configuration file used by your DHCP server is the following command, where *`server-name`* is the name of the executable for your DHCP server:**

```
$ strings `which server-name` | grep dhcpd.conf
```

This will display several strings, one of which will be the full pathname of your DHCP server's configuration file. This works for all of the DHCP servers that I've used—you may find that you need to be slightly more creative with the arguments to the `grep` command, but you'll certainly find your config file with some similar command.

To configure DHCP for your network, you need to add at least the following section (referred to by DHCP config file fans as a stanza) to your DHCP server's configuration file:

```
1: option domain-name "your.domain.name";
2: option domain-name-servers nameserver1-address, nameserver2-address,... ;
3: option routers router-address;
4. option subnet-mask 255.255.255.0;
5: subnet subnet-address netmask netmask-for-that-subnet  {
6:  range full-start-address full-end-address;
7: }
```

NOTE **The actual configuration file used by your DHCP server does not include line numbers (and cannot). I've added the line numbers to make it easier to match up entries in this sample file with the explanations that follow.**

This example provides an extremely minimal DHCP server configuration file that is useful, but which leverages many of the default values provided by the DHCP server. Each line of this stanza has the following meaning:

1. The name of the domain to which any host that receives an IP address from this server should belong. Replace *your.domain.name* with the name of your domain. This configuration file entry is a global setting for the DHCP server, and will apply to all subnet entries that it contains (unless they specifically override this value). On Linux systems, this domain is written to the file `/etc/resolv.conf` without preserving any existing entries. See Bonus Chapter 11, "Setting Up a DNS Server," for more information about the `/etc/resolv.conf` file.
2. The name(s) of any DNS servers that any host that receives an IP address from this server should use to do DNS lookups. Replace *nameserver1-address*, *nameserver2-address*, and so on with the IP addresses of your DNS servers. If you are not running your own DNS server, you can use any DNS servers that your ISP provides, or one of the public DNS servers from sites such as OpenDNS.com (`http://opendns.com/`). This configuration file entry is a global setting for the DHCP server, and will apply to all subnet entries that it contains (unless they specifically override this value). On Linux systems, these name servers are added to the file `/etc/resolv.conf` without preserving any existing entries. See Bonus Chapter 11 for more information about the `/etc/resolv.conf` file.
3. The third line identifies the IP address of the system through which requests for any host that is not found on the local network should be sent. Replace *router-address* with the IP address of your router. This configuration file entry is a global setting for the DHCP server, and will apply to all subnet entries that it contains (unless they specifically override this value).
4. The fourth line specifies the default subnet mask to use on computers that get their IP addresses from this DHCP server. This is typically `255.255.255.0`, but may differ depending upon the configuration of your network. This configuration file entry is a global setting for the DHCP server, and will apply to all subnet entries that it contains (unless they specifically override this value).

NOTE **Each global setting line in the host section ends with a semicolon.**

5. The fifth line identifies the subnet on which the DHCP server should hand out IP addresses. This is expressed by identifying the base address of a network, using zeroes for IP address fields that the DNS server can provide, and by specifying the subnet mask used to derive IP addresses on that network. Replace *subnet-address* with the base address of the network on which you will be handing out IP addresses, and replace *netmask-for-that-subnet* with the appropriate netmask.
6. The sixth line identifies the range of IP addresses that the DHCP server can issue to clients. Replace *full-start-address* with the first IP address that you want your DHCP server to provide, and replace *full-end-address* with the last IP address that you want your DHCP server to provide.
7. The final line terminates the entry for the specified subnet.

After making these changes to your DHCP server's configuration file, you must start (or restart) the DHCP server. If the DHCP server is running and you installed the DHCP server that I recommended, you can kill and restart it by using the following command:

```
$ sudo /etc/rc.d/init.d/dhcp3-server restart
```

All other DHCP servers provide similar start/restart configuration scripts—if you installed another DHCP server, check its control files in `/etc/init.d` for the name of the script that you need to execute and any additional arguments that you may need to supply.

Specifying Additional DHCP Server Configuration File Entries

In addition to the DHCP configuration statements discussed in the previous section, you can specify several other options in your DHCP server's configuration file. Using these other options can provide your DHCP server with additional configuration information, which enables it to forward this information to any of its clients.

Additional global `dhcpd.conf` statements that you may find useful are the following. An example of a `dhcpd.conf` file that uses these statements is given following the list, to clarify the context in which each statement can be used:

- `default-lease-time`: This DHCP server configuration option enables you to specify the default period of time for which a lease will be valid. The time period is specified as an integer number of seconds.
- `ddns-update-style`: This declaration is used if your site uses a DNS (Domain Name Server) server that can be automatically updated as addresses are assigned by the DHCP server and you installed the `autodns-dhcp` package on your DHCP server. If this is the case, see the section of Bonus Chapter 11 entitled "Integrating DHCP with DNS." If your site does not use a DNS server that accepts updates from a DHCP server, you should set this parameter to none to avoid potential performance problems.
- `ip-forwarding`: If, for example, you don't want a DHCP client to be able to forward IP requests to other networks, you may want to set this parameter to `off`. This option is sometimes used on local-only IP networks.
- `max-lease-time`: This DHCP server configuration option enables you to specify the maximum period of time for which a lease will be valid. The time period is specified as an integer number of seconds.

You can also use DHCP to provide specific IP addresses to specific hosts within a subnet group, where a host is identified by MAC address. This is done using the `host` declaration within a subnet group. An example of this, an associated, host-specific command, is the following:

```
authoritative;
option domain-name "vonhagen.org";
option domain-name-servers 192.168.6.1, 208.67.222.222;
option routers 192.168.6.1;
option ddns-update-style interim;
option default-lease-time 28800;
option max-lease-time 86400;
option ip-forwarding off;
subnet 192.168.6.0 netmask 255.255.255.0  {
    range 192.168.6.160  192.168.6.190
    host dorothy {
```

```
            hardware ethernet 00:!E:0C:1D:75:1D;
            fixed_address 192.168.6.69;
            option host_name "dorothy" ;
        }
    }
```

Note that the example for the specific host named `dorothy` overrides any default values specified earlier in the file.

Picking the Right Lease Period

The lease time for your server's DHCP clients is a delicate balance between usability and serviceability. You want clients that have been configured via DHCP to keep that configuration as long as they need that information, but you also want to maximize the availability of DHCP configuration to other clients on your network. This is especially (and primarily) important on networks where the number of potential users exceeds the number of available IP addresses.

Identifying any conceivable patterns in how users access and use your DHCP server will help determine a usable lease period. For example, if you know that the number of possible clients exceeds the number of available IP addresses, it's in your best interest to keep the lease period as short as possible, On the other hand, you want to minimize network traffic and the load on your DHCP server. If you have many more users than IP addresses and mobile users frequently connect to your network, you will want to keep the lease time as short as possible. However, if systems typically appear on your network for a certain period of time (30 minutes, an hour, and so on), you will want to set the lease time to that period to minimize lease refresh requests.

Another consideration is the amount of time required to recover from network problems or hardware failures on the DHCP server. For example, if your DHCP server were to crash and burn, the lease period would define the maximum amount of time that DHCP clients would maintain their current network configuration.

Finally, consider the "silly user" factor. If clients of your DHCP server are setting up network services that they expect others to be able to access for a significant amount of time, you may need to make sure that your DHCP leases are valid for at least that period of time. Even though DHCP leases are supposed to be renewed when they are actively in use, this may not happen for one reason or another. As we all learned in Cub Scouts, Brownies, or similar groups, "Be prepared." If users report problems, you can always fix things and report a success story, but sometimes no comments are the best comments of all.

Managing a DHCP Server Graphically

As you might expect, the Ubuntu repository provides an excellent GNOME utility for graphically configuring a DHCP server. You should have installed this package (`gdhcpd`) as part of the instructions for installing a DHCP server in the section of this chapter entitled "Installing a DHCP Server."

Installing the `gdhcpd` package creates a GDHCPD entry on the Applications ➪ MAInternet menu, which you can select to start Ubuntu's graphical DHCP server configuration tool. Unfortunately, this

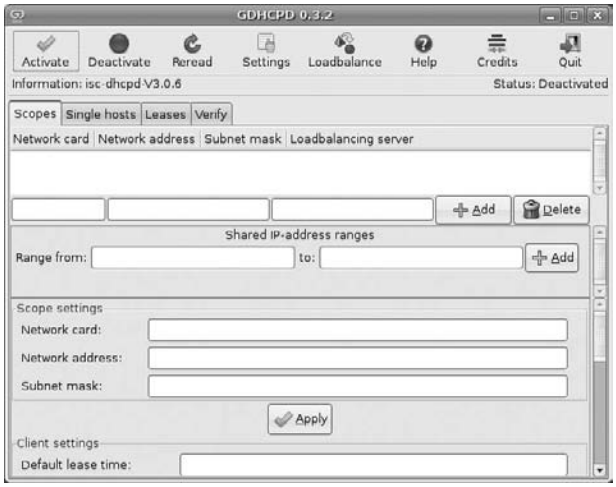
menu item may not start the gdhcpd utility using `sudo`, which would mean that it can't write to the configuration file for your DHCP server (which is the file `/etc/dhcp3/dhcpd.conf` if you installed the ISC v3 DHCP server that I suggested earlier). To eliminate problems related to file protection issues, start the `gdhcpd` command from a GNOME terminal, `xterm`, or other Ubuntu command-line interface using the following command:

```
$ gksudo gdhcpd
```

After you supply your password, the `gdhcpd` utility starts, as shown in Figure B10-1.

FIGURE B10-1

Starting the GNOME DHCP configuration utility



Don't worry if the `gdhcpd` application displays a message about not being able to find the DHCP server in your path because (by default) it looks for a DHCP server executable named `dhcpd`, while the name of the ISC v3 DHCP server that I suggested installing is `dhcp3`.

The first thing to do is to identify the scope for this DHCP server, which is the Ethernet interface that this DHCP will listen on, the possible range of IP addresses that it will deliver, and the subnet mask that it will use. To do this, enter these values in the fields shown in Figure B10-2, which identify a DHCP server running on Ethernet interface `eth0` of the host that will hand out IP addresses on the `192.168.6.0` subnet using a subnet mask of `255.255.255.0`.

NOTE **A *scope* is a way of identifying the complete range of possible IP addresses for a network. A scope typically defines a single physical subnet on your network to which DHCP services are offered. Scopes provide the primary way for a DHCP server to manage distributing and managing a specific range of IP addresses and delivering them (and other network configuration parameters) to clients on that subnet.**

Once you have entered this information, click Apply to define that interface/server/netmask combination. An entry for that server displays on the Scopes tab, as shown in Figure B10-3.

FIGURE B10-2

Identifying the DHCP server's Ethernet information

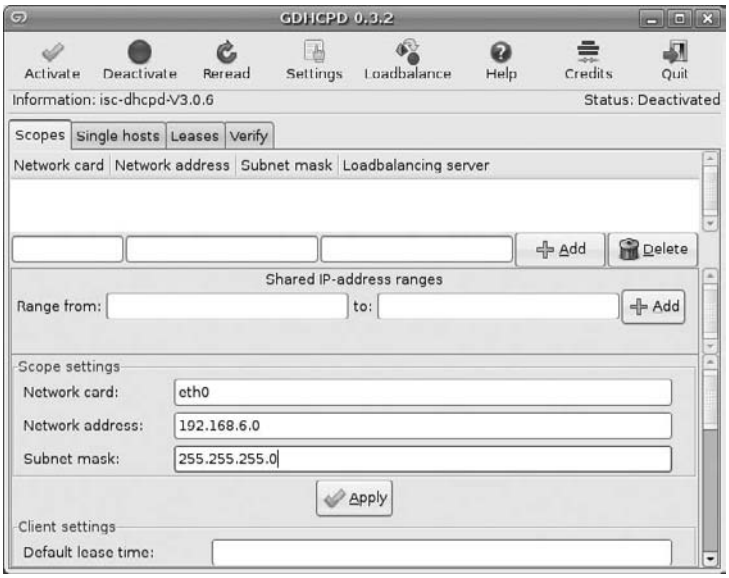


FIGURE B10-3

Successful addition of a DHCP server scope

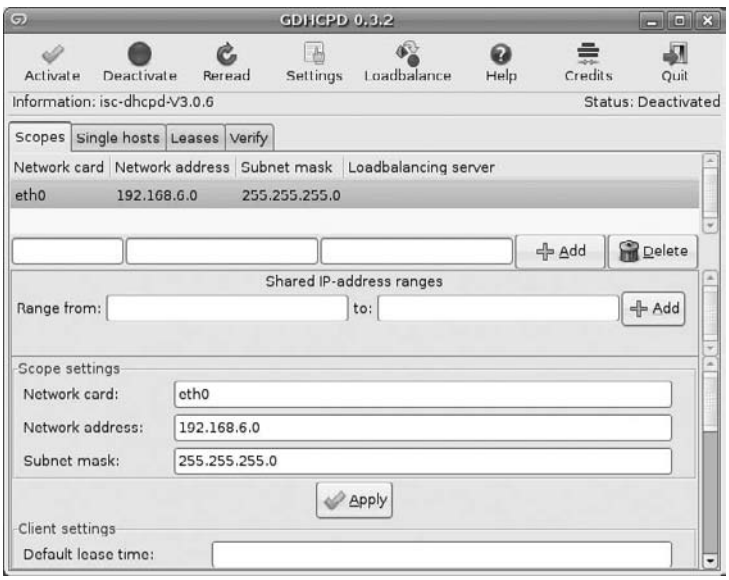
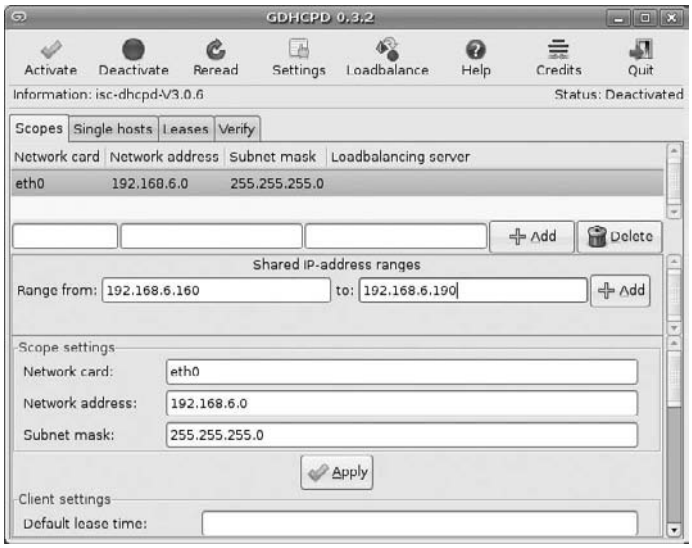


FIGURE B10-4

Identifying a range of DHCP addresses to serve



Next, enter the range of IP addresses that you want this DHCP server to offer to clients, as shown in Figure B10-4.

After entering these values, click Add to define this range. The defined range displays in the primary `gdhcpd` dialog, as shown in Figure B10-5.

Finally, click Settings to examine the default settings associated with your DHCP server. The dialog shown in Figure B10-6 displays.

Change any of the values on this screen that reflect how your DHCP server should behave. If you are configuring a DHCP server other than the dhcp3 server that I suggested installing earlier in this chapter, you may need to change the name of the DHCP server's configuration file and the path to the file in which DHCP leases are stored to reflect the appropriate locations for that DHCP server.

You may also want to change general configuration items. For example, if you are also running a DHCP server and the DHCP server that you are configuring can automatically update that server, you should change the value of the DDNS update style parameter to either `ad-hoc` or `interim` (ad hoc has been deprecated and should not be used; `interim` allows your DHCP server to update a DNS server whenever it hands out a lease). DHCP and DNS server integration also requires some other DHCP configuration file changes; see the section of Bonus Chapter 11 entitled "Integrating DHCP with DNS" for more information.

Once you have verified any global values that you want to set for your DHCP server, click Apply to close this dialog and return to the primary `gdhcpd` dialog.

If you are finished defining ranges of DHCP addresses that you want your DHCP server to offer to clients, click Quit to exit the `gdhcpd` configuration utility. Your DHCP server is configured—all that remains is to start it!

FIGURE B10-5

A defined range of DHCP addresses to serve

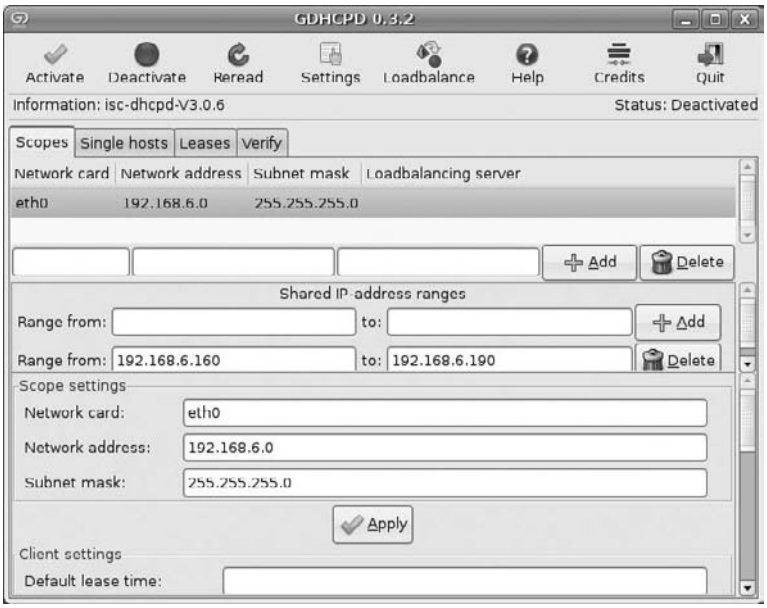
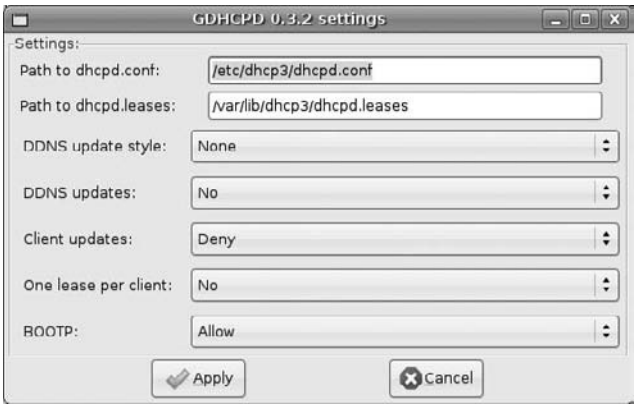


FIGURE B10-6

General DHCP server settings



TIP **Before starting your DHCP server, you may want to check the section of this chapter entitled "Specifying Additional DHCP Server Configuration File Entries" to see if there are any other entries that you want to add to your DHCP server's configuration file.**

To start your DHCP server using the new configuration file, execute the following command:

```
$ sudo /etc/init.d/dhcpd3-server restart
```

This terminates any version of the `dhcpd3-server` application that may be running and forces the process to restart, which reads your new configuration file. Congratulations—you're running a DHCP server.

At this point, you should check the file `/var/log/messages` to verify that your DHCP server started correctly. You can do this using a command like the following, which shows the last 100 lines in this file (just in case):

```
$ sudo tail -100 /var/.log/messages
```

If the DHCP server did not start correctly or clients cannot successfully connect to the DHCP server, see the following section for some debugging and troubleshooting tips.

Troubleshooting DHCP

The most common cause of DHCP problems is that the server did not start correctly or has already handed out all of the leases that are available to it. By default, any startup, status, and error messages from a DHCP server are written to the log file `/var/log/messages`. If your DHCP server does not start correctly, this file is the first place to look. If you have manually created your DHCP server's configuration file, missing semicolons or other syntax errors will cause the DHCP server to exit or, worst-case, start up in some mode other than the one that you intended.

If a client computer is unable to contact a DHCP server, it usually assigns itself an IP address like `169.x.x.x`, which traditionally means that your computer can't connect to the network. If any client system gets an address in this range, the first things to check are the obvious ones—is the client computer actually connected to a port on your local network, and is the Ethernet interface that is connected configured as a DHCP client? On Ubuntu systems, you can check the configuration of your Ethernet interface(s) by selecting the System ➪ Administration ➪ Networking menu item and supplying your password. You should also check the `/var/log/messages` file on your DHCP server to see if the DHCP request from that client was received but rejected for some reason.

On multi-homed systems, a very common problem is that a DHCP server is listening on the wrong Ethernet interface, and is therefore not receiving DHCP client requests on that subnet. Starting a DHCP server on a specific Ethernet interface is done by identifying the name of that Ethernet interface on the command line that starts the DHCP server, which is typically in a file with a name like `/etc/init.d/dhcpd`. The name of this file may be different depending on the DHCP server that you are using.

A final suggestion is to explicitly force the client to release any lease that it may believe it currently has. The `dhclient3 -r` command (or `dhclient -r` command, if you are running a v2 DHCP client) forces the client to explicitly release any leases that it believes that it has. The DHCP client then exits—you can restart it by executing a command like the following:

```
$ sudo /etc/init.d/networking restart
```

This restarts all of the networking-related processes on your system, starting the DHCP client as part of that sequence if your system is indeed configured as a DHCP client.

Summary

Being able to dynamically assign and centrally manage client IP address and network configuration information dramatically simplifies life for many system administrators. This chapter explained how to install a Dynamic Host Configuration Protocol (DHCP) server on your Ubuntu system, using either command-line or graphical software management tools. It then explained how to configure your DHCP server either from the command line or user-friendly graphical tools. The chapter concluded by discussing potential problems that you may encounter when using and configuring DHCP, and provided ways to diagnose and correct them.

Bonus Chapter 11 discusses how to set up your own Domain Name System (DNS) server. A DNS server maps hostnames to IP addresses (and vice versa), and is often a necessity on corporate or academic networks where you must manage your own hostnames and internal IP addresses.

Bonus Chapter 11

Setting Up a DNS Server

As I mentioned in Bonus Chapter 10, a long time ago when even I was young, the Internet (or the ARPANET as it was called in those days) was a small place with a limited number of hosts, all of which had fixed IP addresses. The mappings between those IP addresses and the actual names of those hosts were maintained in a central text file. Yes, that's right, a single text file. Everyone shared this text file (HOSTS.TXT) with everyone else by retrieving it from the Stanford Research Institute's Network Information Center (the one true NIC) every so often, and then adding their own modifications for any private local hosts and networks. However, as the ARPANET grew and more and more hosts became networked, maintaining all of this information in a single file became not only impractical, but silly.

Not only was this file long (and getting longer), but its format was only specified by convention rather than explicit decree. Luckily, the Internet/ARPANET was created by smart people, so many improvements to this model occurred in the early 1980s. For example, in March of 1982, the hosts table definition was updated with DoD Internet Host Table Specification (RFC 810). Shortly thereafter, a server function to provide individual hostname/address translations was described in RFC 811, which introduced the idea of dynamically providing and maintaining host information. RFC 819, "The Domain Naming Convention for Internet User Applications," woke everyone up, and RFC 830, "A Distributed System for Internet Name Service" suggested the one true concept of a distributed system of name servers, each serving its local domain. From 1982 to 1983 was a busy time for the fundamental concepts that we all take for granted today (and may not even be aware of).

TIP **Once the appropriate RFCs came out and people had the chance to think about them, implementation was the next step. By November, 1983, the concept and schedule for actually implementing DNS were developed and published in RFC 881, "The Domain Names Plan and Schedule"; RFC 882, "Domain Names—Concepts and Facilities"; and RFC 883, "Domain Names—Implementation and Specification."**

IN THIS CHAPTER

- Overview of DNS and BIND
- Overview of BIND configuration files
- Creating zone files
- Creating reverse lookup files
- Troubleshooting DNS

In response to these RFCs, the clever folks at the University of California at Berkeley (UCB) created a software package called BIND, for Berkeley Internet Name Daemon. Amusingly enough, this was first written as a graduate student project. BIND was eventually taken over and maintained first by the university's Computer Systems Research Group (CSRG), and then by Digital Equipment Corporation (DEC) for a while, and then eventually by the Internet Systems Consortium (`www.isc.org`), which still maintains BIND today. Many other DNS-related RFCs (973, 1035, and so on) have been published on DNS since the original flurry, some enhancing existing RFCs and some replacing them with newer, more modern and more specific, content. As mentioned in Bonus Chapter 10, "Setting Up a DHCP Server," you can find Internet RFCs at sites such as `www.ietf.org/rfc.html` or the somewhat easier-to-browse site at `www.faqs.org/rfcs`.

NOTE **References to BIND throughout the rest of this chapter should be taken to mean BIND version 9, which is the version that I suggest (and describe) installing and using.**

DNS was primarily designed to address the scalability problems that prompted the migration of hostname/IP address mappings from a central file to a flexible mechanism that can be queried and which returns results. This has some great additional benefits for local subnets in today's ubiquitous networking universe, where private networks can support their own DNS servers that satisfy local, private requests while forwarding requests for external host information to higher-level name servers. These local name servers can also cache retrieved requests, reducing the amount of external network requests your name servers have to issue over time. Many sites run name servers that only serve to cache external requests, and are therefore known as "caching-only name servers."

The remainder of this chapter provides an overview of how DNS works, how to install the latest version of the BIND name server on your Ubuntu or Kubuntu system, how to create the text files required for a basic BIND installation and integrate them into your BIND server, how to restart and test a BIND server, and how to diagnose many common DNS problems.

NOTE **BIND and DNS are rich, complex topics about which many entire books have been written. The last section of this chapter provides pointers to some more detailed, dedicated DNS and BIND resources that you can find online. If you are going to be setting up and managing complex or commercial DNS servers, I'd suggest that you pay a quick visit to any brick-and-mortar or online bookseller for dedicated DNS and BIND reference material that you will always have at your fingertips regardless of the state of your network. Two of my favorites are *DNS and BIND* by Paul Albitz and Cricket Liu (O'Reilly—make sure you get the latest edition!), and *Pro DNS and BIND* by Ron Aitchison (Apress). You can find the latter online at `www.netwidget.net/books/apress/dns`, but you should still buy a copy.**

Overview of DNS and BIND

DNS gives your network an entry point into a hierarchical collection of hostname and IP address servers on an intranet or the Internet itself. All of the hosts on your network can query one or more local DNS servers for local hostname and IP address information. Requests for hostname and IP address information outside the administrative domain of your DNS servers are automatically passed up the DNS food chain until some DNS server can authoritatively give you the information you're looking for.

A common DNS misconception is that a system needs to run a DNS server to participate in DNS. This is not the case. Systems can be pure DNS clients without running a name server process. The client-side portion of the DNS system is known as a *resolver*, which is configured through the file `/etc/resolv.conf` (which I'll discuss later in this section). Whether a system participates in DNS at all is configured

through your system's Name Service Switch file, /etc/nsswitch.conf, which is also discussed in detail later in this section.

Most hosts have traditional sources of hostname/IP address information, such as /etc/hosts files, as well as DNS servers. When an application needs to contact a specific host by name, your system goes through the following process:

1. It first checks the /etc/nsswitch.conf file's hosts entry to determine how your system is configured to resolve hostnames and IP addresses. The /etc/nsswitch (Name Service Switch) identifies the sequence in which files and network resources should be consulted to find a variety of information, including hostname/IP address information. The hosts entry generally looks something like the following:

```
hosts:    files dns [NOTFOUND=return] nis nisplus
```

 The system on which this entry is used first looks in the /etc/hosts file for host information, and then consults DNS. If the host is not found after querying DNS servers, the hostname/IP address lookup fails. Items in a hosts entry after [NOTFOUND=return] are not consulted, but are often left on this line as a reminder that they are valid values for hostname/IP address information. In this example, the NIS and NIS+ information services would not be consulted.

TIP **Some older applications still consult another configuration file called /etc/host.conf. This file provides similar information to that provided in /etc/nssiwtch.conf files about the order in which an /etc/hosts file and DNS should be consulted using the following entry:**

```
order hosts,bind
```

The resolver functions in most recent versions of glibc use /etc/nsswitch.conf. However, if you have changed your /etc/nsswitch.conf file and some specific application is consulting /etc/hosts and DNS in the wrong order, check your /etc/host.conf file. You may also want to report the problem to the current maintainer of the program that you were running.

2. Assuming the hostname or IP address of a specific host is not found in the /etc/hosts file, a system that is configured to consult DNS checks the contents of the text file /etc/resolv.conf to identify the DNS name servers that it should consult, in order. This file also often contains information such as the default domain that should be supplied when looking up non-fully qualified hostnames. For example, the /etc/resolv.conf file on one of my Ubuntu systems looks like the following:

```
domain vonhagen.org
nameserver 0.0.0.0
nameserver 192.168.6.64
nameserver 208.67.222.222
nameserver 208.67.220.22
```

 In this example, non-fully qualified hostnames such as 64bit are assumed to be within the domain vonhagen.org. The name servers are queried in order until one responds or a maximum number of connection failures and associated retries have elapsed, in which case the query fails. The name servers 208.67.222.222 and 208.67.220.22 are public name servers from the folks

at OpenDNS.com, who provide public name servers and related services, and are well worth checking out! For a list of publicly available DNS servers all over the world, see `www.dns-serverlist.org`.

TIP **A name server specification of `0.0.0.0` is actually a shortcut for querying a local name server using the first IP address that was assigned to the local system. This is the loopback address, which is usually `127.0.0.1`, but is not guaranteed to be. Using the shortcut `0.0.0.0` causes your request to be sent to an IP address that is known to be valid for the current host, whatever it may be.**

3. When a DNS server responds, it first tries to resolve the query itself. DNS servers are configured through many specially formatted text files that are discussed in the next section. These files identify any DNS zones, which are essentially analogous to a domain or subdomain that the server can provide authoritative information for. If the hostname or IP address request is in a zone that the DNS server is authoritative for or it is already present in the server's cache of previously answered requests, the DNS server will reply directly to the request. If the request is in a subdomain that the DNS server knows an authoritative source of information for, it will forward the request to that DNS server for resolution. That DNS server will then reply to the client.
4. If the first DNS server that responds cannot resolve the request in one way or another, it queries one of the Internet's top-level name servers (known as *root servers*) to identify the name server that it believes is responsible for the domain or the range of IP addresses that contains the IP address for which the application is looking for hostname information.
5. The root name server returns a pointer to the appropriate name server, which your system then queries for hostname or IP address information. If this name server is not actually authoritative for the domain or IP address that you are querying, it may forward the query to other name servers—this hierarchy continues until a name server is found that is actually authoritative for the domain or IP address that you are querying. This forwarding is generally invisible to you, and most often occurs when querying hosts on subdomains rather than top-level domains.

This combination of local, root, and hierarchical name servers is elegant in its simplicity. When your host finally receives the hostname or IP address that it is looking for, it typically caches that value so that subsequent requests for the same information won't have to go any further than the local name server.

BIND is generally configured in one of four ways:

- As a master server that is an authoritative source of DNS information for one or more zones, which generally map to a domain or subdomain.
- As a slave server that is an authoritative source of DNS information for one or more zones, but periodically retrieves and synchronizes this information from a master server rather being the actual repository for that information.
- As a caching-only server that is not authoritative for any zone, but which caches the results from DNS queries to minimize network traffic and expedite answering multiple requests for the same DNS information. Most Ubuntu server systems will run at least a caching-only name server for these reasons.
- As a resolver-only client that does not run a DNS server but whose DNS resolver is configured to consult name servers through its `/etc/resolv.conf` file.

Most desktop systems do not need to actually run a name server and will work fine as resolver-only clients. However, if you spend a vast amount of time surfing the Web or work on a heavily loaded network, you may want to upgrade your system to running a caching-only client. DNS servers do not consume a vast amount of system resources, but they do require memory, process slots, and so on. If a system that you are configuring or using does not need to provide host information for a zone (which would mandate running a DNS server), only the user(s) of a system can determine which of the caching-only or resolver-only configurations works best for that host.

Installing a DNS Server

Although every Ubuntu and Kubuntu system provides the basic client-side commands for querying and dynamically mapping hostnames to IP addresses, an actual DNS server is not installed as part of a default Ubuntu or Kubuntu installation. Let's face it—not everybody needs to run a DNS server, and for small, home networks, maintaining your `/etc/hosts` files isn't that big of an issue if you have only a few machines. However, if you want to run your own name server (which you will probably want to do if you are responsible for an enterprise network or a SOHO network with a relatively large number of hosts) the package management utilities provided on Ubuntu and Kubuntu systems make it easy to install the BIND DNS server by installing the `bind9` package. See Chapter 23 for more information about installing software on Ubuntu and Kubuntu systems.

TIP **As always with server packages such as BIND, I suggest that you install BIND using `aptitude` from the command line. This will enable you to use the `-r` option to automatically install the packages that are recommended for use with BIND. They don't use much disk space, and having more tools is always better!**

After installing the `bind9` package and friends, setting up your own DNS server is a simple matter of creating a few configuration files, as explained in the next section.

Overview of BIND Configuration Files

The `/etc/nsswitch.conf` and `/etc/resolv.conf` files reflect your system's general approach to resolving hostname and IP address queries. The primary configuration file for the BIND name server, `named.conf`, is found in the directory `/etc/bind` on Ubuntu and Kubuntu systems.

As delivered on Ubuntu and Kubuntu systems, the `/etc/bind/named.conf` file primarily consists of statements that load other BIND configuration files, providing a simple, hierarchical structure for your BIND configuration files. The files that a Ubuntu/Kubuntu BIND server loads fall into several categories:

- Configuration files for the BIND server process, `named`. BIND installations load a general, site-independent configuration file and a separate, site-specific configuration file. This simplifies future updates by making it easier to identify files that may have changed in the BIND installation on a specific system.
- Zone files that define the mapping of hostnames in that DNS zone to IP addresses.
- Reverse lookup files that define the mapping of IP addresses in a numeric DNS zone to hostnames.

As a specific example, the default /etc/bind/named.conf file on a BIND server loads the following files. These are listed in alphabetical order, to simplify finding a file in this list rather than in the order that they are actually loaded:

- /etc/bind/db.0: Contains reverse lookup information that resolves the broadcast addresses 0.*.
- /etc/bind/db.127: Contains reverse lookup information that resolves the loopback addresses 127.*.
- /etc/bind/db.255: Contains reverse lookup information that resolves the broadcast addresses 255.*.
- /etc/bind/db.empty: Contains a shared file that is used for each of the nonroutable IP address families specified by RFC 1918, "Address Allocation for Private Internets" (www.faqs.org/rfcs/rfc1918.html).
- /etc/bind/db.local: Contains hostname to IP address information for the local loopback interface, 127.0.0.1.
- /etc/bind/db.root: Contains hostname to IP address information for the root name servers.
- /etc/bind/named.conf.local: Contains zone entries that identify any zones and IP addresses for which this BIND server is a source of authoritative information. This file is a local, host-specific supplement to the file /etc/named.conf.
- /etc/bind/named.conf.options: Defines the basic options with which the local BIND server has been configured.
- /etc/bind/zones.rfc1918: Provides reverse lookup information for each of the nonroutable IP families specified by RFC 1918, "Address Allocation for Private Internets" (www.faqs.org/rfcs/rfc1918.html). This file is not loaded by Ubuntu/Kubuntu BIND servers by default, but can be loaded in your /etc/bind/named.conf.local file to reduce the time spent trying to resolve IP addresses in nonroutable networks.

In most cases, you will modify only the /etc/bin/named.conf.local file for your BIND server. In this file, you will add entries for files that define the zones and IP address ranges for which your BIND server can provide authoritative DNS information.

TIP **The /etc/bin/named.conf.options file for BIND servers running on Ubuntu and Kubuntu systems configures those servers to use /var/cache/bind as their default working directory. All temporary DNS data is stored in that directory. For this reason, all pathnames in the /etc/bin/named.conf.local file either should be the full pathnames of files located in nontransient directories in your filesystem, such as /etc/bind, or will be interpreted as files that are to be found in /var/cache/bin.**

A sample /etc/bin/named.conf.local file from one of my DNS servers looks like the following:

```
zone "vonhagen.org" {
type master;
file "/etc/bind/vonhagen.zone";
check-names warn;
};

zone "168.192.in-addr.arpa" {
type master;
file "/etc/bind/168.192.in-addr.arpa";
```

```
};

include "/etc/bind/zones.rfc1918.local";
```

As you can see, this DNS server is a master server for the zone `vonhagen.org`, which uses IP addresses in the `192.168.*` family. Information about the zone `vonhagen.org` is provided in the file `/etc/bind/vonhagen.zone`. Reverse lookup information for IP addresses in the `192.168.*` family is provided in the file `/etc/bin/168.192.in-addr.arpa`. As discussed previously, full pathnames must be provided for any files not found in `/var/cache/bind`, which is the default BIND server's default directory. This file also loads a local copy of the default `/etc/bin/zones.rfc1918` file, which I have edited to remove null responses for the zone `"168.192.in-addr.arpa"` which I use locally.

The next section explains how to create the zone and reverse lookup files for any local systems for which your BIND server will provide authoritative information.

Creating DNS Zone and Reverse Lookup Files

As installed on Ubuntu and Kubuntu systems, BIND is a caching-only name server that is appropriate for use on local servers or heavily network-dependent systems to help minimize the amount of external DNS traffic by caching DNS lookups so that subsequent requests for that same information can be provided from the cache rather than by repeating an external DNS request.

Many people want to run their own DNS servers to provide DNS information for the hosts that they are responsible for on their home, enterprise, or academic networks. As discussed earlier in this chapter, DNS information for each set of hosts and associated range of IP addresses is specified in a zone file, one that identifies the IP address associated with a specific hostname (known as `name` to IP address lookups), and one for reverse lookups that identifies the hostname that corresponds to a specific IP address. The next few sections discuss mandatory entries that are commonly used in both types of files, the format and content of a zone file, and the format and content of reverse lookup files.

NOTE **The files that I am calling *zone* and *reverse lookup* share the same format, and are both examples of zone files. I refer to them separately because it's convenient to think of them that way. They contain different types of records and serve different logical purposes in the context of a DNS server.**

Using Common Entries for Zone and Reverse Lookup Files: SOA and $TTL

The beginning of both zone and reverse lookup files contains variable settings for the BIND name server process (`/usr/sbin/named`) and an entry that sets basic variables for each zone or range of IP addresses for which your name server provides authoritative information.

A zone file typically begins with a TTL (Time To Live) entry that provides a default value for any record in the zone file. A TTL value determines how long remote DNS servers can cache any information about that resource that they have retrieved from your DNS server. The default TTL value is preceded with a dollar sign to indicate that you are setting a BIND server variable, and takes a time value expressed in seconds (the default), minutes, hours, or days. The sample TTL value that I generally use is one day (`1d`) because I frequently add, rename, or change hosts on my networks. This setting therefore looks like the following:

```
$TTL 1d
```

The most important entry in a zone or reverse lookup file is the Start Of Authority (SOA) record, which defines global parameters for a zone or range of IP addresses. The format of an SOA record is defined in RFC 1035 (`www.faqs.org/rfcs/rfc1035.html`). A sample SOA record looks like the following:

```
@ 1d IN SOA ubuntu32.vonhagen.org. wvh.vonhagen.org. (
      2008090102 ; Serial
      21600      ; Refresh
      1800       ; Retry
      604800     ; Expiration
      900        ; Negative Cache TTL
)
```

TIP **One confusing thing about the SOA record (and, frankly, about many other DNS configuration file entries) is that many of the fields have default values and are therefore often omitted if the values that you specify are the same as the defaults. Personally, I find the practice of omitting values to be both lazy and the most common cause of DNS problems. The parser for DNS files is quite finicky (to use the technical term), and you will eventually burn yourself or one of your fellow sysadmins by leaving fields out. It's not as if any of the values in an SOA or other records in a zone or reverse lookup file are all that long or require all that much typing. Do everybody a favor—specify all of them.**

The fields in an SOA record are as follows:

- The first field identifies the root name of the zone associated with this SOA record. The @ symbol is a shortcut for this name, which is specified in the zone declaration in the BIND configuration file (in this case, the file `/etc/bind/named.conf.local`). You may occasionally see people put an actual zone name here—I prefer to use the @ symbol because it more clearly differentiates the SOA record from other records in a zone or reverse lookup file.
- The second field is the TTL value to this record. Even though I always put a default value at the beginning of a zone file, I simply repeat that value here to be more precise. This field is technically optional—if it is not specified, the default TTL for the server or an internal default value will be used.
- The third field identifies the record class of the SOA record. An IN entry identifies this as a DNS resource of the "Internet" class. This field is technically optional because IN is the default class type for any resource record in a zone or reverse lookup file, but I prefer to explicitly provide a class value for readability and clarity purposes. Other possible record classes are HS (Hesiod) and CH (Chaos), which are types of records that only MIT staff or Lisp machine fans will recognize.
- The fourth field identifies the name of the name server that provides authoritative information for this zone or IP address range. This name looks like a standard hostname, but usually ends with a period, which is the way in which the DNS configuration file parser identifies fully qualified Internet information. Names that do not end with a period are assumed to be local names to which the root name for the zone or IP range should be appended. For example, if the name of your name server is `ns1`, the root name for this domain is `vonhagen.org`, and the IP address of your name server is identified in the zone or reverse lookup file, then either the entries `ns1` or `ns1.vonhagen.org` are acceptable in this field. However, the entry `ns.vonhagen.org` is completely wrong because it does not end with a period, and the DNS parser therefore replaces it with `ns1.vonhagen.org.vonhagen.org`, which probably doesn't exist in your domain. Forgetting the period here is a common problem when creating zone or reverse lookup files.
- The fifth field is the e-mail address of someone that is responsible for this zone or IP address range. As with other Internet names in a hosts or zone file, this entry usually ends with a period to indicate that it is a fully qualified Internet name. The @ sign that is traditionally used in an

e-mail address is invalid here because the DNS parser uses that as a shortcut for the root name of the zone or reverse lookup range, and is replaced with a period. Again, forgetting the period here is a common problem when creating zone or reverse lookup files. However, it's harder to see this problem because you simply don't get mail sent from the BIND user. If you see several e-mail failures for this user, check the syntax of the SOA records for your BIND server.

- The sixth field is actually a set of five parameters, separated by whitespace and enclosed within one set of parentheses. Many people, including myself, use semicolons (a comment character) to append text descriptions of each field that make it easier to see what each value means during future updates. The following are the five values that you must supply:
 - **serial number:** This is an unsigned integer that should be changed each time you modify a zone or reverse lookup file. This value is usually constructed by concatenating the four digits for the current year, two digits for the current month, two digits for the current day, and two final digits for a version number between 01 and 99. This makes it easy to decipher when the file was last modified and allows up to 99 revisions per day, which should hold just about anyone.
 - **refresh interval:** Identifies the amount of time in seconds after which a slave DNS server will update its zone and reverse lookup information from the master server for that zone.
 - **retry:** Identifies the time interval in seconds after which a slave DNS server whose refresh interval has expired but has failed to contact its master server will try to recontact that master server. This value should obviously be much smaller than the refresh interval.
 - **expiration:** Identifies the time interval between successful refreshes after which the zone data maintained by a slave server is no longer considered authoritative. Slave servers stop responding to DNS requests if this time period has passed and the master server still cannot be contacted.
 - **negative cache TTL:** The amount of time that a negative response, such as a nonexistent domain response, will be cached by the DNS server. Versions of BIND prior to version 9 used this value to represent the default TTL for a zone or reverse lookup file, which is now specified using $TTL (as described earlier).

The next sections describe how to create zone and reverse lookup files that use these common features.

Creating Zone Files

A basic zone file consists of four blocks of information:

- The $TTL and SOA statements described in the previous section.
- An entry that defines the name server for that zone.
- An entry that defines the mail server for that zone.
- Entries for all of the hosts in that zone that map hostnames to their associated IP addresses and that identify any aliases (CNAMEs) that are associated with those hosts.

A name server entry looks like the following:

```
@           . IN   NS    ubuntu32.vonhagen.org.
```

This entry identifies the root name of the zone using the same @ symbol shortcut as was used in the SOA record, the fact that this is an Internet class record (IN), the specific type (NS, for name server), and the hostname of the name server. The IP address of the name server is typically provided later in the zone file

if it is in the same domain or subdomain. As with other Internet name entries in this file, the name of the name server must end with a period or the zone name will be appended to it. I use the @ shortcut to refer to the root domain in NS records because, frankly, this makes it easier for me to use a generic template for all of my zone files.

A mail server entry is very similar to a name server entry except that it uses the MX (Mail Exchanger) record type to identify a mail server, and you can optionally precede the name of the name server with a numeric weight that allows you to identify higher-priority mail servers. Lower values have higher priority. This is meaningless when you have only one mail server, but providing a value makes it easier for you to remember that this feature is available as your network and number of mail servers expand. A sample mail server entry looks like the following:

```
@               IN    MX     10 mail.vonhagen.org.
```

The IP address of the mail server is typically provided later in the zone file if it is in the same domain or subdomain. As with other Internet name entries in this file, the name of the mail server must end with a period or the zone name will be appended to it. I also use the @ shortcut to refer to the root domain in MX records because, again, this makes it easier for me to use a generic template for all of my zone files.

After the name server and mail server records, the rest of a simple zone file typically consists of address (A) and canonical name (CNAME) records for hosts in the zone. Address records look like the following:

```
64bit           IN    A      192.168.6.64
64x2            IN    A      192.168.6.80
ubuntu32        IN    A      192.168.6.90
```

Note that no period follows the short names because I want the root name of the zone to be appended to them. No trailing period is necessary after the IP address.

Canonical name records look like the following, and identify aliases for a host:

```
dualcore        IN    CNAME 64x2
```

No period follows the canonical name or the name of the host that this record provides an alias for. CNAME records require two DNS lookups, one for the name of the alias and a second to find out the IP address of the host that this is an alias for, and so should be used sparingly. CNAME records are most commonly used to map generic server names such as `www`, `ftp`, and so on to a specific machine.

A sample, complete zone file for the domain `vonhagen.org` on one of my local name servers looks like the following:

```
$TTL 1d

@ 1d IN SOA ubuntu32.vonhagen.org. wvh.vonhagen.org. (
       2008090108 ; Serial
       21600      ; Refresh
       1800       ; Retry
       604800     ; Expire
       900        ; Negative Cache TTL
)

@               IN    NS     ubuntu32.vonhagen.org.
@               IN    MX     10 mail.vonhagen.org.
```

```
64bit          IN   A      192.168.6.64
64x2           IN   A      192.168.6.80
ubuntu32       IN   A      192.168.6.90
mail           IN   A      207.44.142.34

dualcore       IN   CNAME 64x2.vonhagen.org.
www            IN   CNAME mail.vonhagen.org.
ftp            IN   CNAME:mail.vonhagen.org.
```

The entries for the `mail`, `www`, and `ftp` hosts are somewhat unique in this file, and merit some additional explanation and an associated apology. Because I use an ISP that already provides a public name server that is authoritative for the domain `vonhagen.org`, I added entries for the IP addresses of specific external hosts to my zone file so that I could find them internally while using this server as an authoritative source for my domain. People on external systems would still use the public authoritative name server for my domain, while internally, I could still find the right hosts for a few external systems. There are other, more complex but standard ways of doing this via delegation, but this hack let me keep my zone file simple while still doing the right thing.

Creating a zone file is half the DNS battle. Once you've created a zone file, you have to create a parallel file that identifies how a name server can deduce the hostname from an IP address, which is known as a reverse lookup file. These files are discussed in the next section.

Creating Reverse Lookup Files

A basic reverse lookup file, sometimes also referred to as a reverse map file, consists of three blocks of information:

- The $TTL and SOA statements described earlier in this chapter.
- An entry that defines the name server for that zone.
- Entries for all of the IP addresses in that zone and the hostnames that they point to.

A name server entry looks exactly the same in a reverse lookup file as it does in a standard zone file:

```
@              . IN   NS     ubuntu32.vonhagen.org.
```

This entry identifies the root name of the zone using the same @ symbol shortcut as was used in the SOA record, the fact that this is an Internet class record (IN), the specific type (NS, for name server), and the hostname of the name server. The IP address of the name server is typically provided later in the zone file if it is in the same domain or subdomain. As with other Internet name entries in this file, the name of the name server must end with a period or the zone name will be appended to it. I use the @ shortcut to refer to the root domain in NS records because, frankly, this makes it easier for me to use a generic template for all of my zone files.

After the name server record, the rest of a simple reverse lookup file typically consists of pointer (PTR) records, each of which identifies the hostname associated with a specific IP address in the range that this file describes. Completely qualified PTR records for IP addresses of the form *AAA.BBB.CCC.DDD* have the following format:

```
DDD.CCC.BBB.AAA.in-addr.arpa. IN  PTR   hostname.domain.tld.
```

The entries in the first field that correspond to an IP address are backward from the actual IP address because this file tells how to map IP addresses back to hostnames, and you therefore want to start with the least significant portion of the IP address, which most easily and uniquely identifies a specific PTR record. All PTR records must include the extension `.in-addr.arpa.`, either explicitly or by having it appended to the entry in the first field.

Sample pointer records from one of my reverse lookup files look like the following:

```
64.6    IN  PTR    64bit.vonhagen.org.
80.6    IN  PTR    64x2.vonhagen.org.
90.6    IN  PTR    ubuntu32.vonhagen.org.
```

Although the entries in the first field represent portions of an IP address, these entries are handled just like names, in the sense that because they are not terminated with a period, the root name of the zone (in this case, `168.192.in-addr.arpa`) is appended to them. The hostnames are terminated with a period to identify them as fully qualified domain names.

A sample, complete reverse lookup file for IP addresses in the family below `192.168`, which are defined in the file `168.192.in-addr.arpa` on one of my local name servers, looks like the following:

```
$TTL 1d

@ 1d IN SOA ubuntu32.vonhagen.org. wvh.vonhagen.org. (
      2008090108 ; Serial
      21600      ; Refresh
      1800       ; Retry
      604800     ; Expire
      900        ; Negative Cache TTL
)

@        IN  NS     ubuntu32.vonhagen.org.

64.6     IN  PTR    64bit.vonhagen.org.
80.6     IN  PTR    64x2.vonhagen.org.
90.6     IN  PTR    ubuntu32.vonhagen.org.
```

Creating a reverse lookup file is the second half of the DNS battle. All that remains is to incorporate them into the BIND configuration for your name server, as described in the next section.

Incorporating Zone and Reverse Lookup Files with BIND

The previous two sections explained how to create zone and reverse lookup files. The files created in the previous two sections are `vonhagen.zone`, the zone file for the domain `vonhagen.org`, and the reverse lookup file `168.192.in-addr.arpa`, which maps IP addresses under `192.168` to the hostnames in the `vonhagen.org` domain that they are associated with. To actually use the files you've created, you simply need to add them to the `/etc/bin/named.conf.local` file for the name server that you want to serve as an authoritative source for this information. The `/etc/bin/named.conf.local` file on this server looks like the following:

```
zone "vonhagen.org" {
type master;
file "/etc/bind/vonhagen.zone";
check-names warn;
};
```

```
zone "168.192.in-addr.arpa" {
type master;
file "/etc/bind/168.192.in-addr.arpa";
};

include "/etc/bind/zones.rfc1918.local";
```

These entries include each of these files and identify this server as the master server for the zone `vonhagen.org`, which uses IP addresses in the `192.168.*` family. Information about the zone `vonhagen.org` is provided in the file `/etc/bind/vonhagen.zone`. Reverse lookup information for IP addresses in the `192.168.*` family is provided in the file `/etc/bin/168.192.in-addr.arpa`. The entries that identify these files use full pathnames because these files are stored in `/etc/bind`, not the server's default directory of `/var/cache/bind` as defined in the server's `/etc/bin/named.options` file. As mentioned previously, this `/etc/bind/named.conf.local` file also loads a local copy of the default `/etc/bin/zones.rfc1918` file, which I have edited to remove null responses for the zone `"168.192.in-addr.arpa"`, because valid values for this range of nonroutable IP addresses are explicitly provided in the file `/etc/bin/168.192.in-addr.arpa`.

After updating your `/etc/bin/named.conf.local` file, all that remains is to reload this information into the name server and test the name server, as described in the next section.

Restarting and Testing Your Name Server

The DNS server startup file for BIND 9 is the file `/etc/init.d/bind9`. Although this startup script provides reload and force-reload options, these options use a name server control utility (`rndc`) that may not be correctly configured (yet) on your system. For this reason, I tend to simply restart the BIND daemon using a familiar command like the following:

```
$ sudo /etc/init.d/bind9 restart
```

After supplying your password, you will see the familiar shutdown and startup messages from the script. You should then verify that the name server started correctly by checking for it in a process listing, using a command like the following:

```
$ ps alxww | grep named
```

If your configuration and zone files are valid, you will see a named process running on your system. Congratulations—you're running a DNS server! If not, proceed to the next section for some debugging tips, and then return to this section to complete your testing.

Next, add an entry for your new DNS server to the file `/etc/resolv.conf`. On the system where you are actually running the DNS server, you can add an entry like the following as the first name server entry in that system's `/etc/resolv.conf` file:

```
nameserver 0.0.0.0
```

As mentioned previously, a name server specification of `0.0.0.0` is just a shortcut for querying a local name server using the first IP address that was assigned to the local system. This is the loopback address, which is usually `127.0.0.1`, but is not guaranteed to be. Using the shortcut `0.0.0.0` causes your request to be sent to an IP address that is known to be valid for the current host, whatever it may be.

Once you're sure that your name server is working correctly, you add the actual IP address of the host on which the name server is running to the `/etc/resolv.conf` file on all of your other systems.

Now it's time to test the DNS server. You can use several commands for testing a DNS server. The most common of these are the following:

- `dig`: The recommended tool for DNS testing nowadays, its output takes a bit of getting used to because of the richness of the information that it returns.
- `host`: The simplest tool for initial DNS testing.
- `nslookup`: Officially deprecated, this command is still found everywhere and is quite useful. Unfortunately, one of these days, the `nslookup` command will stop displaying a warning that it is deprecated, and will simply go away.

All of these commands are easy to use to test your new name server. The following is some sample output from each of these commands on the sample name server that I've configured in this chapter:

```
$ host 64bit
64bit.vonhagen.org has address 192.168.6.64

$ nslookup 64bit
Server:         127.0.0.1
Address:        127.0.0.1#53

Name:   64bit.vonhagen.org
Address: 192.168.6.64


$ dig 64bit.vonhagen.org
; <<>> DiG 9.3.2 <<>> 64bit.vonhagen.org
;; global options:  printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 47870
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 1, ADDITIONAL: 1

 ;; QUESTION SECTION:
;64bit.vonhagen.org.         IN     A

 ;; ANSWER SECTION:
64bit.vonhagen.org.   86400  IN     A      192.168.6.64

 ;; AUTHORITY SECTION:
vonhagen.org.         86400  IN     NS     ubuntu32.vonhagen.org.

 ;; ADDITIONAL SECTION:
ubuntu32.vonhagen.org. 86400 IN     A      192.168.6.90

 ;; Query time: 1 msec
;; SERVER: 127.0.0.1#53(127.0.0.1)
;; WHEN: Wed Dec  5 01:18:21 2008
;; MSG SIZE  rcvd: 91
```

The output of the `host` and `nslookup` commands is straightforward. The output of the `dig` command is the most verbose and useful for debugging purposes because it actually identifies the query that was sent to DNS, identifies the answer that it received, and lists the name server that was identified as being authoritative for the domain in which a host is located. For complete information on the `dig` command and its many options, see the online reference information for this command (`man dig`).

Troubleshooting DNS

One nice feature of BIND 9 is that the named process will not start if there are fatal errors in any of your configuration files. This can be quite irritating until you resolve them, but it does help prevent you from accidentally starting a mangled BIND server. Also, the definition of a fatal error is somewhat flexible. Syntax and parser errors that prevent all of your configuration files from being read correctly are fatal errors, while others that seem equally important, such as missing files or missing or mangled records in a configuration file, are not.

If your BIND server does not start correctly, the first place to look for information is the system log file on your Ubuntu system, `/var/log/syslog`. By default, BIND logs a variety of events to this file when you start or stop the named process, when events such as file reloads occur while the named process is running, and so on. A sample section of the `/var/log/syslog` file showing BIND messages is as follows (I've removed the timestamp and host identification information that would normally begin each line):

```
named[19897]: starting BIND 9.3.2 -u bind
named[19897]: found 1 CPU, using 1 worker thread
named[19897]: loading configuration from '/etc/bind/named.conf'
named[19897]: /etc/bind/named.conf:60: missing ';' before end of file
named[19897]: loading configuration: failure
named[19897]: exiting (due to fatal error)
```

In this case, you can see that there is a missing semicolon on line 60 of my `named.conf` file. After correcting one error in a configuration file, you usually want to check the configuration file again, to make sure that you've eliminated all of the errors, not just one. The Ubuntu `bind` package includes a utility called `named-checkconf` that does just this. Running this utility provides an easy way of checking the syntax of a BIND configuration file without restarting the named process over and over, checking the log each time, fixing one problem or the other, and so on.

The following is another common message that you may see in `/var/log/syslog`:

```
named[20222]: zone vonhagen.org/IN: loading master file: file not found
```

Oddly enough, the fact that the zone file `vonhagen.zone` could not be found is not a fatal error, so the BIND server is running—it just isn't providing any name information for the `vonhagen.org` domain. Aside from syntax errors in the filename, the most common cause of this error message is forgetting to provide the full pathname for an included file. As mentioned previously, BIND on Ubuntu and Kubuntu systems is configured to use the directory `/var/cache/bind` as its working directory. It will look for any files without full pathnames in this directory instead of looking in `/etc/bind`.

The `named-checkconf` file checks the syntax of your `/etc/bind/named.conf` file and other configuration files that it includes—it does not validate the contents or syntax of your zone files. Luckily, Ubuntu's `bind` package includes another utility, `named-checkzone`, that does just that. Although it provides several options, the most common way of using the `named-checkzone` utility is to simply

provide its two mandatory arguments: the name of the zone that a file describes, and the name of the file that describes that zone. This can be either a standard zone file or a reverse lookup file because both of these are different types of zone files that have different logical purposes and contents. Some sample error output from running this utility follows:

```
$ named-checkzone vonhagen.org vonhagen.zone
dns_rdata_fromtext: vonhagen.zone:8: near '@': extra input text
zone vonhagen.org/IN: loading master file vonhagen.zone: extra input text

 # named-checkzone 168.192.in-addr.arpa  168.192.in-addr.arpa
168.192.in-addr.arpa:11: unknown RR type 'ubuntu32.vonhagen.org.'
zone 168.192.in-addr.arpa/IN: loading master file: unknown class/type
```

Don't worry—I mangled copies of these files to illustrate common errors! The ones shown earlier in this chapter are still valid.

The first error, "extra input text," indicates a parser error when trying to read one of the records in this file—in this case, the SOA record. To generate this error message, I removed the closing parenthesis from the SOA record, which causes the parser to continue on to the next line and report that it found extra, unexpected text in the next line, which began with an @ symbol.

The second errors, "unknown RR type 'ubuntu32.vonhagen.org'" and "unknown class/type," are somewhat more insidious. These also indicate parser errors in reading one of the resource records, in this case on line 11. This is the name server declaration in my reverse lookup file. As I mentioned earlier in this chapter, the DNS file parser is very finicky. Although it is clever about using default values whenever possible, you still have to be very careful about how you use whitespace in the file. In this case, I had removed the leading @ symbol and deleted all of the whitespace at the beginning of this record, so that the IN declaration was at the beginning of the line. Because it appeared at the beginning of the lines, without any whitespace preceding it, this was interpreted as the root name of my name server entry, the NS entry was used as the record class, and the name of my name server, `ubuntu32.vonhagen.org` was interpreted as a record type specification. Attempting to resolve this problem also displayed the message "zone 168.192.in-addr.arpa/IN: has no NS records" in the file `/var/log/syslog`, which was caused by a similar parser error that I stumbled over en route to a valid name server entry.

Trying to resolve parser errors in zone files can be frustrating. Worst case, you can fall back to providing an explicit value for every possible field, verifying that the file works correctly, and then simplifying the file until it looks the way that you want it to. Examples, such as those provided earlier in this chapter, are always useful for the purpose of comparison.

Parser and syntax errors can prevent your configuration files from being read at all or from being interpreted correctly. However, once these files are read and interpreted correctly, they may still not work correctly. The most common problem with any zone file is forgetting to append a period to what you intended to be a fully qualified domain name, or putting a superfluous period at the end of a name that you meant to be concatenated with the root name for a zone. Names that do not end with a period are assumed to be local names to which the root name for the zone or IP range should be appended. For example, if the name of your name server is `ns1`, the root name for this domain is `vonhagen.org`, and the IP address of your name server is identified in the zone or reverse lookup file, then either the entries `ns1` or `ns1.vonhagen.org`. are acceptable as the name of the name server. The first is a short name that does not end with a period, which means that the root name of the zone will be appended to it. The second is a fully qualified domain name that ends with a period, so nothing will be appended to it. However, the entry `ns.vonhagen.org` is completely wrong because it does not end with a period, and

the DNS parser therefore replaces it with `ns1.vonhagen.org.vonhagen.org`, which probably doesn't exist in your domain.

This section has presented some of my "favorite" errors in BIND configuration and zone files, and some ways to identify and resolve them. You will undoubtedly see other error messages—unfortunately, this is part of the baptismal process of configuring and starting a BIND DNS server. Be prepared: First, don't panic, and second, have a browser open with your favorite search engine preloaded.

TIP **If you are running a DNS server on a host whose IP address is assigned by a DHCP server (usually through a reserved MAC Address to IP Address mapping), the `/etc/dhcp3/dhclient-script` will overwrite your existing `resolv.conf` entries with the list that is maintained by the DHCP server each time that it checks or renews your lease. The best solution is to add your new DNS server to the official list maintained by the DHCP server. If you don't want to do this for some reason, you can add a line like the following to the file `/etc/dhcp3/dhclient.conf`:**

```
prepend domain-name-servers 0.0.0.0;
```

This would put the name for the local host's DNS server as the first entry in your `/etc/resolv.config` file. Alternately, you can simply comment out the `make_resolv_conf()` function in the script file `/etc/dhcp3/dhclient-script` to prevent your system's `/etc/resolv.conf` file from being updated in the first place, but this cancels out some of the dynamic aspects of DHCP.

Getting More Information About DNS and BIND

As mentioned previously, DNS and BIND are rich, complex topics about which many entire books have been written. This chapter provided an overview of DNS and BIND and explained simple BIND configuration and troubleshooting. For additional information, see one of the books mentioned earlier or consult one of the following online resources:

- **BIND 9 Administrator Reference Manual (`www.net.cmu.edu/groups/netdev/docs/bind9/Bv9ARM.html`):** This is a copy of the official BIND 9 Administrator's manual, done by the BIND folks.
- **Pro DNS and BIND (`www.netwidget.net/books/apress/dns`):** This is an online copy of one of the DNS/BIND books mentioned earlier in this chapter.
- **DNS Resources Directory (`www.dns.net/dnsrd`):** This is a directory site that provides links to a huge collection of other BIND resources.

If you'll pardon the expression, these resources can help you resolve almost any DNS problem.

Summary

This chapter explained how a name server maps hostnames to IP addresses in a flexible, distributed fashion, and explained how to set up your own name server for the hosts on your local network, while still leveraging the other zillion name servers on the Internet. In real life, you would want to set up more than one name server on a local network of any significant size, to protect against the failure of any single system, but this chapter focused on setting up a single name server for smaller networks.

Bonus Chapter 12 continues our exploration of important system servers by explaining how to install and configure a print server. Unless you have more money and surface area than most people, you will want to attach printers to a small number of systems on your network and point your other systems to those print servers for printing purposes. Not only does this save money and simplify administration, but it also helps minimize the number of places that you have to explore to find missing printouts.

Bonus Chapter 12

Setting Up a Print Server

If you have a single system and a single printer, setting up and configuring printing is quite straightforward, and was explained in the section "Adding a Printer" in Chapter 24. However, in today's more complex networked environments, the chances are that you want to access a printer on one system from many other systems, including machines that may run operating systems other than Linux for some legacy software or game-playing reason.

This chapter explains how to set up and tweak connectivity from other computer systems so that your printer is available to everyone else. For your convenience, this chapter also highlights the steps necessary to connect to print servers on both Ubuntu and Kubuntu systems from Microsoft Windows and Mac OS X systems. It concludes by providing some troubleshooting tips and discussions of common problems, as well as additional sources of information.

IN THIS CHAPTER

Enabling remote access to a print server

Integrating Microsoft Windows systems

Integrating Apple Mac OS X systems

Troubleshooting

TIP **Hacking the Ubuntu printing system's configuration files and using unauthenticated printing as described in this chapter is really suitable only for home, SOHO, or SMB environments that are firewalled from the outside world, and in which you hopefully trust everyone. If you have Microsoft Windows systems in your network environment, you may simply want to set up a Samba Server on the system to which your printer is attached. This will leverage your existing Windows authentication mechanisms (depending on how you configure Samba), and will also work with Apple's Mac OS X systems because of its built-in support for Windows printing and SMB. Setting up a Samba Server is discussed in Bonus Chapter 14, "Setting Up a Samba Server."**

Linux and UNIX Printing History

Today's printers are typically high-quality laser or inkjet printers, often capable of color printing and near photographic quality. The original UNIX printing system, known as *Line Printer Daemon (LPD)* was designed to queue and print jobs that were intended for huge, text-only line printers. As more sophisticated

printers were developed that were capable of higher-quality printouts (such as the original Xerox x9700, Canon-CX, and Imagen-300 laser printers), the original LPD print system continued to be used, but required that the jobs that you were printing be preprocessed so that they contained the special commands that the printer used internally to produce higher-quality printouts. This quickly became tedious because it required users to know to which printer they wanted to print to use the appropriate preformatting commands. Eventually, the LPD system was updated and a similar printing system known as LP was developed, which encapsulated the knowledge about the formats required by specific printers. LP implemented the necessary preformatting commands internally by automatically executing them as filters (also known as print drivers) that performed the right formatting and other printer-specific commands before sending the jobs to the target printer. Other updated printing systems, such as `lprng` (lpr, Next Generation, based on the name of the print command used by the LPD system) have also flourished—both LPD and `lprng` are available in the Ubuntu repositories if you need them for compatibility reasons (or because they are the devil that you already know).

The evolution of multiple printing systems for UNIX systems led to incompatibilities between the different print systems, requiring recompilation of the filters for specific printers on different UNIX systems (if you could get the source code at all) and so on. Eventually, a company known as Easy Software Products began developing a more generalized printing system for UNIX, Linux, and other UNIX-like systems known as Common UNIX Printing System (CUPS). The original version of CUPS used the standard networked LPD protocol, but quickly switched to using a new standard, the Internet Printing Protocol (IPP), which non-UNIX/Linux systems such as Windows can use to print to CUPS printers. Easy Software Products also had the foresight to make the CUPS source code freely available under the GPL so that it could be compiled for multiple operating systems and thus become a true, cross-system standard popularized by zillions of users and system administrators. And this strategy has worked—today, CUPS is used in every major Linux distribution and most other UNIX-like systems.

Most Linux distributions, including Ubuntu, provide their own utilities for setting up printers and doing some basic configuration of the printing subsystem. Ubuntu uses the GNOME printer configuration tool, which was described in the section "Adding a Printer" in Chapter 24. However, in addition to standardizing how printing works across many Linux systems, CUPS added one especially nice gift for system administrators: a built-in administrative tool for CUPS configuration that is easily accessed through any Web browser via port 631. In this chapter, however, I am focusing on using the standard Ubuntu/GNOME printer configuration tools as much as possible.

Enabling Remote Hosts to Access Your CUPS Print Server

As distributed with most Linux distributions (including Ubuntu releases prior to 6.06), CUPS servers listen for incoming requests on port 631 of your host's external IP address. With Ubuntu 6.06 and later distributions, the CUPS server only listens for requests on your host's loopback interface, 127.0.0.1, which effectively eliminates any hope of network printing. Luckily, this is easy enough to correct.

NOTE **To avoid having to put the complete documentation for the CUPS configuration file, the print server described in this chapter is a simple print server without any special authentication requirements. This may be unsuitable for WAN enterprise or academic environments.**

To enable network printing on your Ubuntu or Kubuntu system, you must remove this restriction. You can do this by editing the printer's configuration file `/etc/cups/cupsd.conf` manually, or by using

the graphical configuration applets that are provided on both Ubuntu and Kubuntu systems. These are the same utilities that you use to add a printer, select a default printer, and so on—you're just using additional menus to do more advanced administrative tasks.

Ironically, Ubuntu systems (after version 6.06) and Kubuntu 8.10 (and later) systems use the same printer configuration utility, while Kubuntu 8.04 and earlier Kubuntu systems (i.e., all KDE 3-based Kubuntu systems), use an older, KDE-specific printer configuration utility. These two mechanisms for printer configuration are discussed in the next two sections.

TIP **After enabling remote hosts to print to the printers that are provided by your print server, you may still need to configure a printer on remote Linux systems. Whether or not this is necessary depends on the capabilities of the remote Linux system. Most newer Linux distributions support some form of automatic printer discovery. If you can't see the printers on your print server from another Linux system, you will need to configure the printer manually.**

When manually configuring remote printing from another Linux system to your Ubuntu print server, you need to use an Internet Printing Protocol URI of the following form:

```
ipp://hostname/printers/queuename
```

For example, if I were setting up a connection to the CUPS printer named SuperScript-660i running on the host `ubuntu32.vonhagen.org`**, I would supply the URI** `ipp://ubuntu32.vonhagen.org/printers/SuperScript-660i`**. Depending on the type of other Linux systems that you are using, you may have to specify the IP address and queue name separately. In that case, the queue name would be** `printers/queuename`**.**

Enabling Remote Printing on Ubuntu and KDE 4-Based Systems

On Ubuntu systems, start the printer configuration applet by selecting the System ➪ Administration ➪ Printing menu command. On a Kubuntu 8.10 or newer system, select the K Menu ➪ Applications ➪ System ➪ Printing menu command. A dialog like the one shown in Figure B12-1 displays. This dialog is from an Ubuntu system, but the dialog will look essentially the same on Kubuntu systems, with the exception of the window decorations.

This dialog initially shows the printers that you have specifically configured on your system. Select the Server ➪ Settings menu command to display the dialog shown in Figure B12-2.

FIGURE B12-1

Starting the printer configuration applet

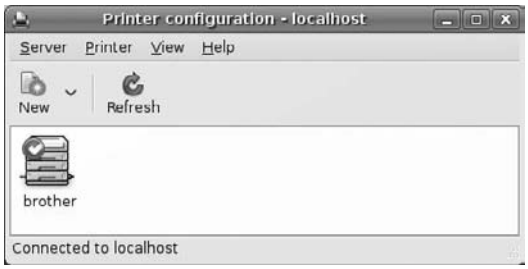


FIGURE B12-2

Configuring printer settings



To enable other hosts to see the printers that are exported by this print server, click the Publish shared printers connected to this system checkbox. To enable other hosts to actually print to these printers, select the Allow printing from the Internet checkbox that is available once you activated the Publish shared printers… checkbox.

To enable this system to see printers that are shared by other print servers on your network, click the Show printers shared by other systems checkbox. You will certainly want to do this if you are sharing printers that are directly connected to any Microsoft Windows or Mac OS X systems that are also on your network.

TIP **If you want to be able to administer this print server from other hosts on your network, select the Allow remote administration checkbox on this dialog. I'd generally recommend against this unless you trust all users with administrative privileges on their systems (which is at least one user on each of your Ubuntu or Kubuntu systems) to administer, configure, and reconfigure your print servers.**

Click OK to proceed. To ensure that the new settings are picked up by the CUPS daemon that is running as the print server on your system, restart the CUPS daemon on your system using the following command:

```
$ sudo /etc/init.d/cupsys restart
```

At this point, other systems on your local network should be able to see any printers that are attached to your system when browsing the list of available printers. If you selected the Show printers shared by other systems checkbox, you should be able to see other printers on your local network, as shown in Figure B12-3.

To verify that other systems on your network can connect to your print server, you can connect to its Web interface using the actual IP address of your system. As mentioned earlier, the default behavior of an Ubuntu 6.06 and later or Kubuntu 8.10 and later system is to accept only connections from your system's loopback IP address, which is the host `localhost` and the IP address `127.0.0.1`. To demonstrate that other IP addresses can contact your print server, you can try to contact its Web interface using your system's actual (i.e., non-loopback) IP address. If the CUPS server is indeed listening on that port on your print server's external Ethernet interface, you will see a page like the one shown in Figure B12-4.

FIGURE B12-3

The printer applet showing printers shared by other systems

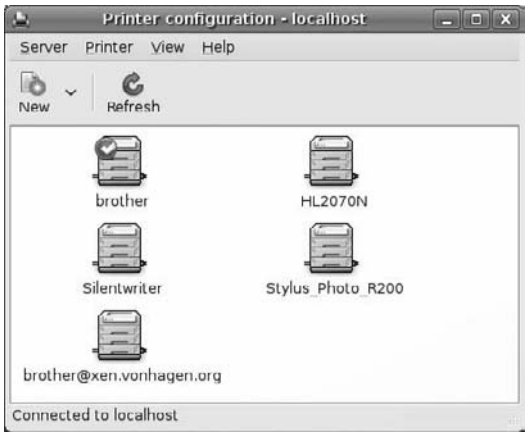
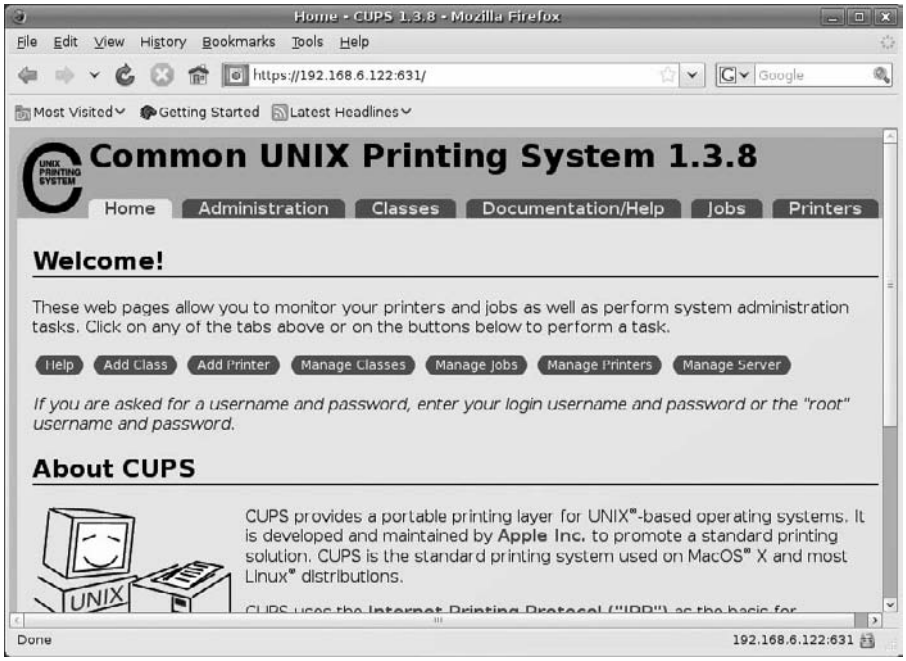


FIGURE B12-4

The CUPS server's Web interface



You should now be able to go to any of the other hosts on your network and create printer entries there. Adding a printer on an Ubuntu system was discussed in Chapter 24. If you are creating network printer entries for this printer on Microsoft Windows or Apple Mac OS X systems, see the subsequent sections of this chapter for tips on how to define network printers on these operating systems.

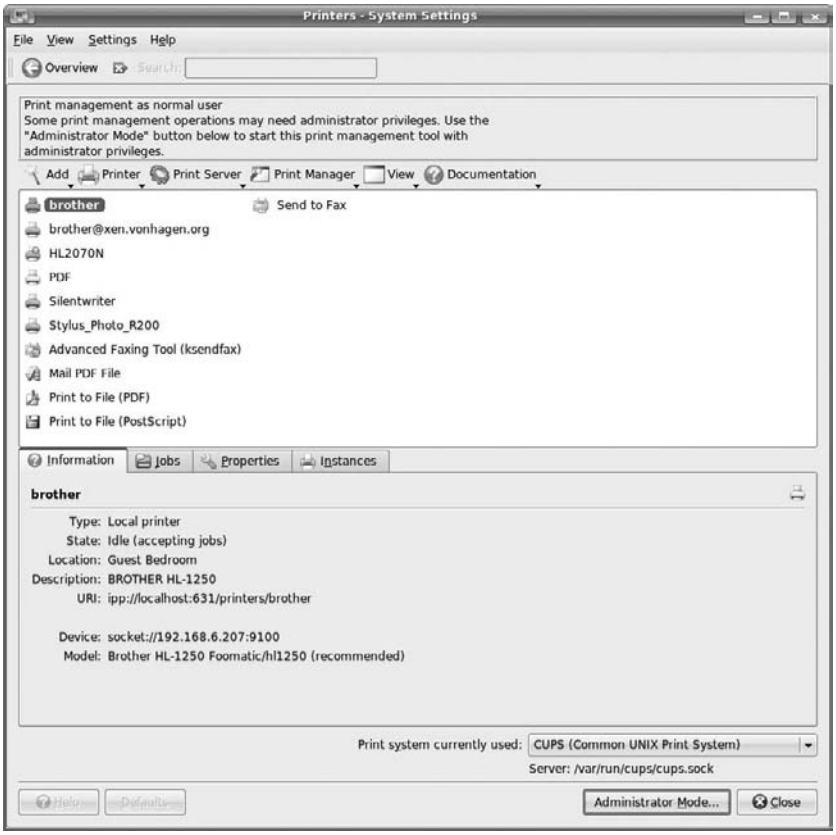
Enabling Remote Printing on KDE 3-Based Systems

KDE 3-based Kubuntu systems (Kubuntu 8.04 and earlier) use a standard KDE-based configuration utility to administer and configure print servers. In most cases, you won't have to specifically enable remote access to your print server on KDE 3-based systems, but this section explains the relevant settings, just in case.

To enable remote printing on a KDE 3-based Kubuntu system (or to verify that remote hosts should be able to see the printers that are supported by this print server), select the K Menu ➪ System Settings menu item and click the Printers icon in the screen that displays. This starts the KDE Printers applet, which displays the dialog shown in Figure B12-5.

FIGURE B12-5

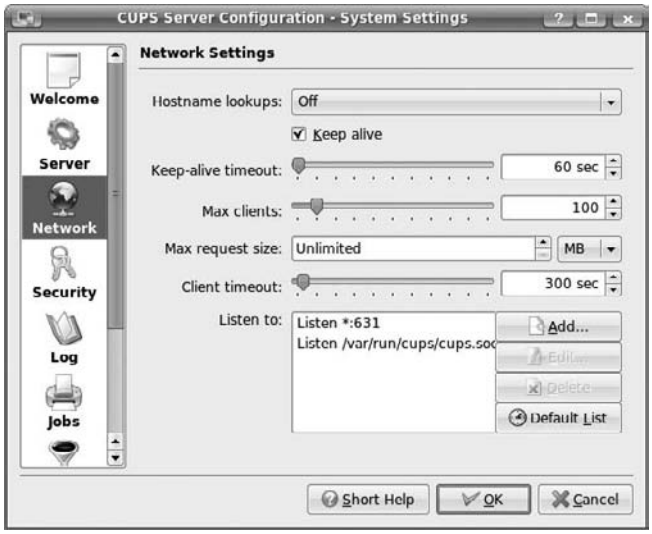
The KDE 3 printers applet



Once this dialog displays, select the Print Server ➪ Configure Server menu command. You will see a warning screen about unknown configuration options, which simply means that the applet doesn't recognize all of the configuration options for the latest CUPS print server—but luckily, it won't remove them or change them. Click OK to close this dialog and display the standard KDE CUPS Server Configuration dialog's Welcome panel. Click the Network icon in the left panel to modify the hosts that can access your printer and modify other network-related configuration settings. This displays the dialog shown in Figure B12-6.

FIGURE B12-6

Configuring network settings in the KDE 3 printers applet



If the `Listen to:` text box contains a setting like `Listen *:631`, this means that your printers should already be available to other systems because the CUPS server is listening on all available network interfaces. If the `Listen to:` text box contains a setting like `Listen localhost:631`, you will need to either remove this entry or replace `localhost` with * so that your CUPS daemon will also listen on your system's external Ethernet interface. To do this, select the Listen `localhost:631` entry, click Edit, specify * as the address to listen on, and click OK to save the modified entry. If, for some reason, you want to prevent other systems from printing on this system's printers, you can specify `Listen localhost:631` as the port for your print server to listen to.

Once you have made any changes that you need to make on this screen, click OK to save those changes and return to the dialog shown in Figure B12-5. Select the Print Server ➪ Restart Server menu command to restart the print server with the new settings.

TIP **If you are having problems administering a KDE 3-based CUPS server remotely, you may need to select the Print Manager ➪ Configure Print Manager menu entry, and click the CUPS Server icon in the left pane to view or reset the identity of the user who is authorized to remotely administer the print server.**

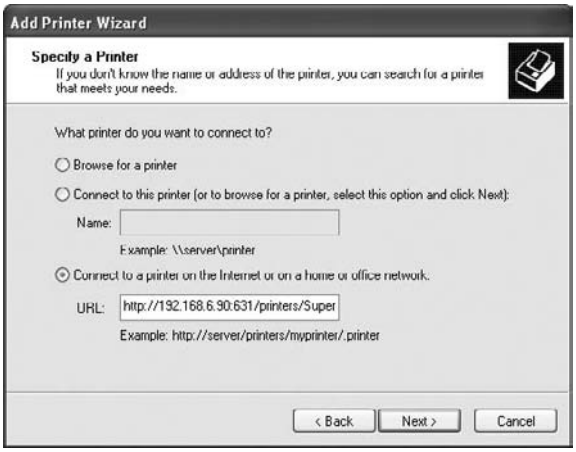
Integrating Windows Printing with an Ubuntu Print Server

Configuring a Microsoft Windows 2000 or XP system to print to a remote Ubuntu printer is as simple as setting up access to any other shared network printer. The following example provides screenshots from a sample Windows XP system, but the process is very similar on earlier versions of Microsoft Windows:

1. Select Start ➪ Settings ➪ Printers and Faxes.
2. Click the Add a Printer task. A hand-waving dialog displays that does nothing. Click Next to proceed.
3. Specify that you want to create a remote printer, and click Next. The dialog shown in Figure B12-7 displays.

FIGURE B12-7

Providing the URL of your Ubuntu or Kubuntu printer



4. Enter a URL of the form `http://host:631/printers/queuename`. For example, to define a connection to a remote printer named SuperScript-660i that is attached to the host `ubuntu32.vonhagen.org`, I would enter the URL `http://ubuntu32.vonhagen.org:631/printers/SuperScript-660i`. Click Next to proceed. The dialog shown in Figure B12-8 displays.
5. Select the manufacturer and model of your printer so that Windows knows how to format documents before sending them to the printer. If the manufacturer or model of the remote printer isn't available in this dialog, you can either select a similar or compatible model or locate a Windows

driver for your printer on the Web or on the CD that fell out of the box when you unpacked your printer. Click OK once you've made your selection.

6. A final hand-waving dialog displays. Click Finish to close that dialog and return to the Printers and Faxes dialog.

FIGURE B12-8

Identifying your printer's manufacturer and model



That's all there is to it! If you want to verify that your printer works correctly, right-click on the icon for the new printer, select Properties, and click Print Test Page. A standard Microsoft Windows printer test page should emerge from your printer shortly.

Integrating Mac OS X Printing with an Ubuntu Print Server

Apple's Mac OS X has come a long way from previous versions of the Mac OS. Mac OS X is a true UNIX-like system that provides an easy-to-use interface. Because Mac OS X includes and uses a huge collection of open source software, integrating your Mac OS X system with a remote Ubuntu print server is easily done with the familiar Mac OS X Printer Setup and Printer Browser tools. The screen shots in this section were taken from a sample Mac OS X 10.4 system, but the process is very similar on other versions of Mac OS X.

To create a Mac OS X definition for your remote Ubuntu printer, do the following:

1. Start the Printer Browser. You can do this several ways. I typically go to the Applications folder in a Finder window, select the Utilities folder, double-click the Printer Setup Utility, and click Add in its menu bar to start the Printer Browser. Regardless of how you get there, Figure B12-9 shows the Printer Browser with no printers defined.

FIGURE B12-9

Starting the OS X 10.4 Printer Browser



TIP **If you are using Mac OS X 10.3 or earlier, hold down either Option key while clicking Add, and skip to Step 3 of this procedure to end up at the equivalent dialog.**

2. Hold down either Option key and click the More Printers button.
3. Click the drop-down menu at the top of this dialog and select Advanced, as shown in Figure B12-10.
4. Click the Device drop-down menu and select Internet Printing Protocol using IPP. This displays a dialog like the one shown in Figure B12-11.
5. Enter anything that you'd like as the Device Name, and enter the Device URI (Universal Resource Indicator) for your Ubuntu printer in the following form:

```
ipp://host:631/printers/queuename
```

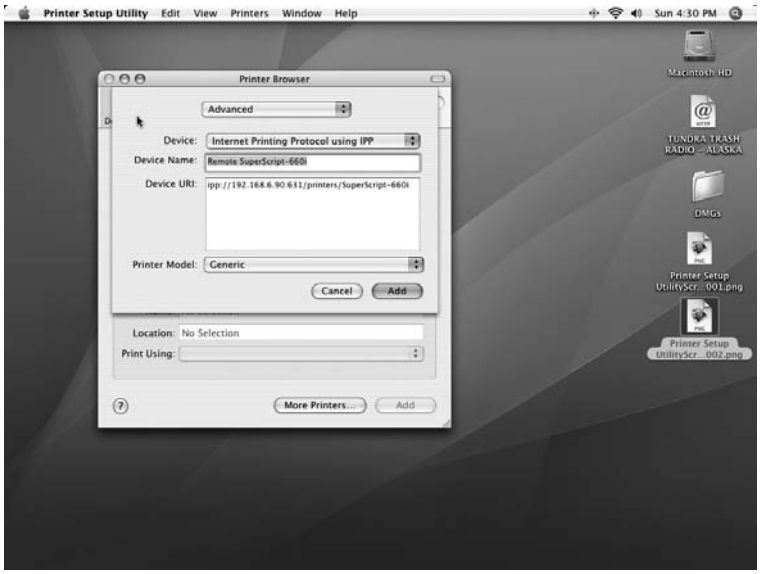
FIGURE B12-10

Specifying advanced printer configuration



FIGURE B12-11

Specifying a name and the URI for your remote printer



For example, to define a connection to a remote printer named SuperScript-660i that is attached to the host `ubuntu32.vonhagen.org`, I would enter the URI `http://ubuntu32.vonhagen.org:631/printers/SuperScript-660i`.

6. If you are using anything other than a generic PostScript printer, click Printer Model, and select the manufacturer and model of your printer. Figure B12-12 shows a driver being selected for an HP LaserJet 4–compatible printer.

FIGURE B12-12

Identifying the manufacturer and model of your printer



7. Click Add to proceed.

That's all there is to it! Your Ubuntu printer is now available for use from your Mac OS X system. Because Mac OS X doesn't provide an easy way of generating a printer test page, you can fire up your favorite application and print a sample page from there. Your print job should emerge from your printer shortly.

Troubleshooting Remote Printing

This section summarizes a few common sources of diagnostic information and potential problems that you may need to try, depending on the version of CUPS that you are running and whether or not you have typed every URI correctly.

Checking the CUPS Log Files

CUPS print servers maintain three log files that provide a variety of information about attempts to access or use them. These log files, stored in the directory `/var/log/cups`, are as follows:

- `access_log`: Contains log messages for each remote host that attempts to access the CUPS print server.
- `error_log`: Contains log messages that record any errors encountered or produced by the CUPS print server.
- `page_log`: Contains log messages for every page printed by a CUPS server. These log messages also summarize information about the host from which the print job was received, the name of the printer being used, and so on.

Of these, the `access_log` and `error_log` files are the most useful for diagnostic purposes. Examining the end of these files after encountering a print failure usually shows useful, meaningful error messages. The `access_log` file will identify the IP address of the systems that have connected to your print server. This can help you verify that the remote system is actually connecting to the print server. Once you've established this, check the `error_log` file for tips about why the job may not have printed.

Accessing Controls for Portions of the CUPS Web Interface

The default CUPS configuration on Ubuntu systems tends to change with different releases, reflecting both enhancements and current security concerns. It's almost certain that the CUPS print server will be updated at some point while you are using this and subsequent releases of Ubuntu. If remote printing works fine but then ceases to work, you may need to add your remote hosts (or your entire network) to the list of acceptable locations in the CUPS daemon's primary configuration file, `/etc/cups/cupsd.conf`. The list of valid locations for incoming print jobs is stored inside the `<Location /> </Location>` stanza. If printing suddenly stops working, make sure that this stanza contains an entry that either enables access by default or explicitly allows access from the printer from all hosts on the local network. For example, the following is a restrictive version of this stanza that denies access from anywhere other than `localhost` or any hosts on the `192.168.6.*` network:

```
<Location />
  Order Deny,Allow
  Deny From All
  Allow From localhost
  Allow From 192.168.6.*
</Location>
```

The CUPS configuration file that is provided as part of a standard Ubuntu installation at the time of this writing allows access from anywhere (once you've changed CUPS to actually listen on your print server's external Ethernet interface, as described earlier in this chapter). However, changes to the default configuration file are common, and it's pretty depressing to have printing cease to function after a system "update." For example, different portions of the CUPS administrative interface each have their own `<Location>` entries. *If* you cannot access them from your system, try adding an `Allow` entry for a specific host (rather than from an entire subnet, as described previously).

Similarly, most administrative functions on modern CUPS servers require that you contact them using SSL (i.e., a URI that begins with https). You can either disable this requirement by commenting out the lines in your `cupsd.conf` file that contain the word Encryption, or use SSL. When using SSL, you may

need to create an exception for the default SSL certificate used on Ubuntu and Kubuntu systems because this is a self-signed certificate that is not guaranteed by any recognized certificate authority. You will also have to do this on every system that you want to use to contact the CUPS server's Web-based interface. Your browser will walk you through the steps to create an exception for your print server's certificate.

Handling Preformatted Print Jobs

Another common update-related problem may arise if you receive an update that changes how CUPS handles different types of incoming print jobs. If you are printing from a remote system, such as a Microsoft Windows or Apple Mac OS X system, that system preformats the print jobs before sending them to the print server. If printing suddenly stops after an update and you see error messages in the `error_log` file about enabling raw input, make sure that the following entry is present in the file `/etc/cups/mime.types` and is not preceded by a hash mark:

```
application/octet-stream
```

The hash mark is a comment character in CUPS configuration files. This entry must be present for preformatted print jobs to be received via HTTP and sent to the printer.

After making any changes to any CUPS configuration file, you need to restart the CUPS daemon on your system using the following command:

```
$ sudo /etc/init.d/cupsys restart
```

Getting More Information About CUPS

As you have probably guessed, CUPS configuration is either extremely rich and powerful or extremely irritating—and usually both. The CUPS configuration file `/etc/cups/cupsd.conf` has so many possible options that it makes even the command-line `gcc`, `ls`, and `tar` commands jealous. If you try to do anything serious with CUPS beyond the simple print server described in this chapter, you will want more information. The following are some excellent sources of additional information about CUPS:

- `www.cups.org/documentation.php`: The online documentation for CUPS that is available from Easy Software, Inc.
- `www.linuxprinting.org`: Specializing in help with printing on GNU/Linux systems, this site features a very useful database of printers that are known to work with Linux (`www.linuxprinting.org/printer_list.cgi`) and a good deal of useful information about CUPS, including a great CUPS troubleshooting tutorial. (I left out the URL for the latter intentionally because it would line-wrap into the next chapter.)
- *CUPS: Common UNIX Printing System* (Michael R. Sweet, SAMS, 2002; ISBN: 0672321963): Michael Sweet is one of the owners of Easy Software, Inc., and thus, this is a pretty definitive book, not just for CUPS, but also for the Internet printing protocol. Unfortunately, the book is a bit old now, but it is still an excellent resource and my copy is a valuable asset.

As with other Ubuntu topics, you should also be able to get a significant amount of useful information from the online Ubuntu Forums discussed in Chapter 1 of this book.

Summary

Centralizing print services on one or a small number of computer systems is a good idea for almost any network. This chapter provided a bit of background about printing on UNIX and Linux systems, and then discussed how to install and configure CUPS, the Common UNIX Printing System. This chapter also provided specific information about enabling any Microsoft Windows or Mac OS X systems that you may still have lying around to print to your Ubuntu or Kubuntu print server.

Bonus Chapter 13 explains how to share directories over a network using NFS, the Network File System, which is the standard distributed file system mechanism for all UNIX and Linux systems. Sharing a central repository of project files, documents, and code over a network in a work environment is important from many standpoints, including increasing efficiency, minimizing confusion, and minimizing the size of your backups. Today, it is just as important in many home environments, enabling everyone on your home network to share access to your household's music collection, vacation photos, and any other files that you want your house mates, spouse, or children to be able to see.

Bonus Chapter 13

Setting Up an NFS Server

IN THIS CHAPTER

- Overview of NFS
- Installing an NFS server
- Configuring an NFS server
- Mounting NFS filesystems
- Debugging NFS problems

Sharing groups of files that multiple people need access to is standard operating procedure in business today and, thanks to home networking, is getting to be Standard Operatig Procedure for home use as well. Providing centralized access to a collection of audio materials that you've extracted from your CD collection or the vacation photos from your most recent trips is just as important to the home user as providing centralized access to your procedure manuals and software source repository is to the business user or SOHO developer. Luckily, Linux systems provide several ways of sharing directories over a network, some oriented primarily toward Linux and other UNIX-like systems, including Apple's Mac OS X, and others oriented more toward Microsoft Windows systems (which Linux and Mac OS X systems can also access, of course). This chapter discusses how to set up one of your Ubuntu Linux systems so that other systems can access its directories over the network using NFS, the Network File System, which is popularly used on all Linux and UNIX-like systems. (For information on setting up your Ubuntu system to share directories with Microsoft Windows systems, see Bonus Chapter 14, "Setting Up a Samba Server.")

Sun Microsystems' Network File System, better known simply as NFS, is the most common networked filesystem in use on UNIX and UNIX-like systems today, largely because it is available and free for every UNIX and UNIX-like system I've seen since 1990 or so (including Mac OS X), and is also available for Microsoft Windows. Microsoft Windows networked filesystems that use CIFS (SMB) may be a close second—a sad, but unavoidable, business reality. However, because NFS was designed for systems whose filenames are case-sensitive and understand UNIX-style permissions, NFS is still a superior choice for UNIX-like systems such as Linux.

This chapter provides an overview of NFS, discusses different versions of NFS and their capabilities, and discusses the various applications associated with NFS. Beyond this background material, this chapter focuses on explaining how to set up your Ubuntu system to be an NFS file server—accessing NFS file servers from other systems was explained in Bonus Chapter 2. NFS is commonly

used with Kerberos, a networked authentication protocol that provides secure authentication in any networked environment by exchanging encrypted keys. Setting up Kerberos authentication is discussed in Bonus Chapter 4.

Overview of the Network File System

NFS is a network filesystem that provides transparent access to files residing on remote disks. Network filesystems are often commonly referred to as *distributed filesystems* because the files and directories that they provide access to may be physically located on many different computer systems that are distributed throughout your home, academic environment, or business. Developed at Sun Microsystems in the early 1980s, the NFS protocol has been revised and enhanced several times between then and now, and is available on all Linux, UNIX, and UNIX-like systems (including Mac OS X), and even for Windows systems from many third-party software vendors. The specifications for NFS have been publicly available since shortly after it was first released, making NFS a de facto standard for distributed filesystems.

NFS is the most common distributed filesystem in use today, largely because it is free and available for almost every type of modern computer system. NFS enables file servers to export centralized sets of files and directories to multiple client systems. Good examples of files and directories that you may want to store in a centralized location but make simultaneously available to multiple computer systems are users' home directories, site-wide sets of software development tools, and centralized data resources such as mail queues and the directories used to store Internet news bulletin boards.

The following are some common scenarios for using NFS:

- **Sharing common sets of data files:** Sharing files that everyone on your network wants to access, whether they are audio files, business data, or the source code for tomorrow's killer app, is the most common use of any type of networked filesystem.
- **Explicitly sharing home directories:** Suppose that the home directories for all of your users are stored in the directory `/export/home` on your NFS file server, which is automatically mounted as `/home` on all of your UNIX and UNIX-like systems at boot time. Users can then log in on any NFS client system and instantly see their home directory, which would be transparently made available to them over the network. If you want to have a local `/home` directory for some users and an NFS-mounted home directory for others, you could mount `/export/home` from your NFS server elsewhere (perhaps `/home/remote` or something similarly evocative) and manually configure specific users to have home directories located in this other directory.

TIP **An alternative to the previous bullet is to automatically mount networked home directories using an exported NFS directory that is managed by an NFS automount daemon known as `autofs` on Linux systems. Whenever access to a directory managed by an automount daemon is requested by a client, the daemon automatically mounts that directory on the client system. Automounting simplifies the contents of your server's `/etc/exports` file by enabling you to export only the parent directory of all home directories on the server, and letting the automounter manage that directory (and therefore its subdirectories) on each client. See the sidebar later in this chapter for general information on automounting, a complete discussion of which is outside the scope of this chapter.**

- **Sharing specific sets of binaries across systems:** Suppose that you want to make a specific set of GNU tools available on all of the systems in your computing environment, but also want to centralize them on an NFS server for ease of maintenance and updating. To ensure that configuration files were portable across all of your systems, you might want to make these binaries available

in the directory `/usr/gnu` regardless of the type of system that you were using. You could simply build binaries for each type of system that you support, configuring them to be found as `/usr/gnu` but actually storing them in directories with names such as `/export/gnu/ubuntu`, `/export/gnu/solaris8`, and so on. You would then configure each client of a specified type to mount the appropriate exported directory for that system type as `/usr/gnu`. For example, `/export/gnu/ubuntu` would be mounted as `/usr/gnu` on Ubuntu systems, `/export/gnu/solaris8` would be mounted as `/usr/gnu` on Solaris systems, and so on. You could then simply put `/usr/gnu/bin` in your path and the legendary "right thing" would happen regardless of the type of system that you logged in on.

As you'll see in this chapter, NFS is easy to install and easy to configure, and provides a flexible networked filesystem that any Ubuntu, other Linux, UNIX, or UNIX-like system can quickly and easily access. In some cases, it's easy to trip over a few administrative gotchas, but Ubuntu provides powerful and easy-to-use tools that simplify configuring NFS file servers to "do the right thing."

How NFS Works

If you simply want to use NFS and aren't too concerned about what's going on under the hood, you can skip this section. However, this section provides the details of many internal NFS operations because some enquiring minds do indeed want to know and because, frankly, it's just plain interesting to see some of the hoops that NFS clients and servers have to use to successfully communicate between different types of computer systems, often with different types of processors. So, if you're interested, read on, McDuff!

The underlying network communication method used by NFS is known as *Remote Procedure Calls* (*RPCs*), which can use either the lower level Universal Datagram Protocol (UDP) as their network transport mechanism (NFS version 2) or TCP (NFS version 3 and later). For this reason, both UDP and TCP entries for port 2049, the port used by the NFS daemon, are present in the Linux `/etc/services` file. UDP minimizes transmission delays because it does not attempt to do sequencing or flow control, and does not provide delivery guarantees—it simply sends packets to a specific port on a given host, where some other process is waiting for input.

The design and implementation of RPCs make NFS platform-independent, interoperable between different computer systems, and easily ported to many computing architectures and operating systems. RPCs are a client/server communication method that involves issuing RPC calls with various parameters on client systems, which are actually executed on the server. The client doesn't need to know whether the procedure call is being executed locally or remotely—it receives the results of an RPC in exactly the same way that it would receive the results of a local procedure call.

The way in which RPCs are implemented is extremely clever. RPCs work by using a technique known as *marshalling*, which essentially means packaging up all of the arguments to the remote procedure call on the client into a mutually agreed-upon format. This mutually agreed-upon format is known as eXternal Data Representation (XDR), and provides a sort of computer Esperanto that enables systems with different architectures and byte-orders to safely exchange data with each other. The client's RPC subsystem then ships the resulting, system-independent packet to the appropriate server. The server's RPC subsystem receives the packet, and unmarshalls it to extract the arguments to the procedure call in its native format. The RPC subsystem executes the procedure call locally, marshalls the results into a return packet, and sends this packet back to the client. When this packet is received by the client, its RPC subsystem unmarshalls the packet and sends the results to the program that invoked the RPC, returning this data in exactly the same fashion as any local procedure call. Marshalling and unmarshalling, plus the use of the

common XDR data representation, make it possible for different types of systems to transparently communicate and execute functions on each other.

RPC communications are used for all NFS-related communications, including communications related to the authentication services used by NFS (NIS or NIS+), file locks, NFS mount requests, status information, and requests made to the NFS `automount` daemon. To enable applications to contact so many different services without requiring that each communicate through a specific, well-known port, NFS lets those services dynamically bind to any available port as long as they register with its central coordination service, the `portmapper` daemon. The `portmapper` always runs on port 111 of any host that supports RPC communications, and serves as an electronic version of directory assistance. Servers register RPC-related services with the `portmapper`, identifying the port that the service is actually listening on. Clients then contact the `portmapper` at its well-known port to determine the port that is actually being used by the service that they are looking for.

Communication failures occur with any networked communication mechanism, and RPCs are no exception. As mentioned at the beginning of this section, UDP does not provide delivery guarantees or packet sequencing. Therefore, when a response to an RPC call is not received within a specific period of time, systems will resend RPC packets. This introduces the possibility that a remote system may execute a specific function twice, based on the same input data. Because this can happen, all NFS operations are *idempotent*, which means that they can be executed any number of times and still return the same result—an NFS operation cannot change any of the data that it depends upon. Even though NFS versions 3 and later use TCP as their network transport mechanism, the idea of idempotent requests is still part of the NFS protocol to guarantee compatibility with NFS version 2 implementations.

As another way of dealing with potential communication and system failures, NFS servers are stateless, meaning that they do not retain information about each other across system restarts. If a server crashes while a client is attempting to make an RPC to it, the client continues to retry the RPC until the server comes back up or until the number of retries exceeds its configured limit, at which time the operation aborts. Stateless operation makes the NFS protocol much simpler because it does not have to worry about maintaining consistency between client and server data. The client is always right, even after rebooting, because it does not maintain any data at that point.

Although stateless operation simplifies things, it is also extremely noisy, inefficient, and slow. When data from a client is saved back to a server, the server must write it synchronously, not returning control to the client until all of the data has been saved to the server's disk. As described in the next section, "Comparing Different Versions of NFS," newer versions of NFS do some limited write caching on clients to return control to the client applications as quickly as possible. This caching is done by the client's `rpciod` process (RPC IO daemon), which stores pending writes to NFS servers in the hopes that it can bundle groups of them together and thus optimize the client's use of the network. In the current standard version of NFS (NFS version 4), cached client writes are still essentially dangerous because they are only stored in memory, and will therefore be lost if the client crashes before the write completes.

In a totally stateless environment, a server crash would make it difficult to save data that was being modified on a client back to the server once it is available again. The server would have no way of knowing what file the modified data belonged to because it had no persistent information about its clients. To resolve the problem, NFS clients obtain file handles from a server whenever they open a file. *File handles* are data structures that identify both the server and the file that they are associated with. If a server crashes, clients retry their write operations until the server is available again or their timeout periods are exceeded. If the server comes back up in time, it receives the modified data and the file handle from the client, and can use the file handle to figure out which file the modified data should be written to.

NFS 4 provides client-side caching, which can significantly improve performance by reducing the number of calls that need to be made to an NFS 4 server in order to retrieve data, and which further helps prevent data loss if a server is unavailable for a short period of time. However, this is not sufficient to insulate NFS clients from performance problems if they must execute binaries located on an NFS server as part of the client's reboot process. If the server is unavailable, the client cannot boot. For this reason, most NFS clients must contain a full set of system binaries, and typically only share user-oriented binaries and data via NFS.

Comparing Different Versions of NFS

NFS has been around almost since the beginning of UNIX workstation time, appearing on early Sun Microsystems workstations in the early 1980s. This section provides an overview of the differences between the four different versions of NFS, both for historical reasons, and to illustrate that NFS is by no means a done deal. NFS 4 resolves the biggest limitations of NFS 3, most notably adding real client-side data caching that survives reboots. The most common versions of NFS used on systems today are NFS versions 4 and 3, which are the versions that this chapter focuses on. You will probably want to use NFS 4 whenever possible, but may still need to use NFS 3

The following list identifies the four versions of NFS and highlights the primary features of each:

- **Version 1:** The original NFS protocol specification was used only internally at Sun during the development of NFS, and I have never been able to find any documentation on the original specification. This would only be of historical interest.
- **Version 2:** NFS version 2 was the first version of the NFS protocol that was released for public consumption. Version 2 used UDP exclusively as its transport mechanism, and defined the 18 basic RPCs that made up the original public NFS protocol. Version 2 was a 32-bit implementation of the protocol, and therefore imposed a maximum file size limitation of 2GB on files in NFS and used a 32-byte file handle. NFS version 2 also limited data transfer sizes to 8KB.
- **Version 3:** NFS version 3 addressed many of the shortcomings and ambiguities present in the NFS version 2 specification, and took advantage of many of the technological advances in the 10+ years between the version 2 and 3 specifications. Version 3 added TCP as a network transport mechanism, making it the default if both the client and server support it; increased the maximum data transfer size between client and server to 64KB; and was a full 64-bit implementation, thereby effectively removing file size limitations. All of these were made possible by improvements in networking technology and system architecture after NFS version 2 was released. Version 3 also added a few new RPCs to those in the original version 2 specification, and removed two that had never been used (or implemented in any NFS version that I've ever seen). To improve performance by decreasing network traffic, version 3 introduced the notion of bundling writes from the client to the server, and also automatically returned file attributes with each RPC call, rather than requiring a separate request for this information as version 2 NFS had done.
- **Version 4:** Much of the NFS version 4 protocol is designed to position NFS for use in Internet and World Wide Web environments by increasing persistence, performance, and security. Version 4 adds persistent, client-side caching to reduce network traffic and decrease the chance of data loss in a reboot, and adds support for ACLs and extended file attributes in NFS filesystems. Version 4 also adds an improved, standard API for increased security through a general RPC security mechanism known as Remote Procedure Call Security - Generic Security Services (RPCSEC_GSS). This mandates the use of the Generic Security Services Application Programming Interface (GSS-API, specified in RFC 2203) to select between available security mechanisms provided by clients and servers.

The version of the NFS server that is shipped with Ubuntu provides support for NFS versions 4, 3, and 2.

NFS Users and Authentication

NFS uses the user ID (UID) and group ID (GID) of each user from a system's password file (`/etc/passwd`) to determine who can write to and access exported files and directories, based on the UID and GID that own those directories on the file server. This means that all of your users should have the same UID and GID on all systems to which NFS directories such as home directories are exported.

In small networks, it is often sufficient to make sure that you create the same user and groups on all of your systems, or to make sure that the password and group files on your file server contain the correct entries for all of the user and groups who will access any directory that it exports.

In larger networks, this is impractical, so you will want to consider network-oriented authentication mechanisms, such as Kerberos, which is discussed in Bonus Chapter 4, or the traditional Network Information System (NIS), which was developed by Sun Microsystems specifically for use with NFS. Unfortunately, discussing NIS installation and setup on Linux systems is usually a history lesson rather than something that modern sysadmins actually do, so it is therefore outside of the scope of this chapter. If you need to use NIS on Linux, you can find a variety of excellent information about it online in documents such as the NIS HOWTO (`www.linux-nis.org/nis-howto/`), which is even available in multiple languages. NIS was great in its day and can still be quite useful in networks of heterogeneous and aging UNIX systems, but it's not the way to go in a modern networked environment comprising Linux systems, Mac OS X systems, and the occasional Microsoft Windows system.

Installing an NFS Server

You can use the Synaptic Package Manager, the Adept Package Manager, or the command-line package management utilities described in Chapter 23 to install the packages required to run and monitor an NFS server on your Ubuntu or Kubuntu system. You must install both the `nfs-common` and `nfs-kernel-server` packages on any system where you want to run an NFS server—only the `nfs-common` package is required on a system on which you only want to mount NFS filesystems. (Installing the `nfs-kernel-server` package on a server will automatically install the `nfs-common` package as a requirement.) Once the installation completes, you're ready to configure your server to share data from your system with any system that supports NFS.

TIP **If you are going to want to use Kerberos as the authentication mechanism for your NFS server (and your computing environment in general), see the section of Bonus Chapter 4 entitled "Configuring Kerberbos Authentication." This chapter focuses on setting up an NFS server in small networks where common user and group IDs can be manually maintained, although Kerberos requirements or alternatives are explained wherever possible.**

Configuring an NFS Server

Older releases of Ubuntu provided a graphical tool called `shares-admin` that could be used on NFS servers that had a GUI to simplify the process of setting up the directories that were to be exported from that NFS server. This tool is no longer supported and, although it is still delivered with Ubuntu Hardy 8.04,

should not be used—it does not configure export entries for NFS 4 correctly. This section therefore focuses on configuring an NFS server from the command line because that method works on all Ubuntu desktop, Ubuntu server, and Kubuntu systems.

As described in the next few sections, configuring an NFS server is a relatively simple process, consisting of setting some parameters in two configuration files that identify the services that should be started, and defining the directories that you want to export in a text file called `/etc/exports`.

Given that the NFS kernel server in the Ubuntu repositories supports NFS versions 2, 3, and 4, you should usually configure a new NFS server as an NFS 4 server, as is done in the next few sections. As long as the `/etc/exports` file explicitly defines all exported directories, NFS 4 servers can still respond to mount requests from NFS clients that support only NFS 2 or NFS 3. The primary differences between configuring an NFS 4 server and configuring NFS 2 or 3 servers are:

- NFS 4 servers export directories using a pseudo-filesystem, which is just a directory that you create with specific protections on your NFS 4 server and define specially in your `/etc/exports` file.
- NFS 4 adds new configuration options for the exported directories that are listed in the `/etc/exports` file.

The next few sections explain each of these steps, highlighting differences between NFS servers that expect common UID and GID values, and those that use a networked authentication mechanism such as Kerberos.

NOTE **The parameter settings in the NFS configuration files discussed in the next two sections have no values by default. An NFS server will attempt to infer the correct settings for these when it starts. However, to ensure compatibility with mount requests from clients that can only speak older versions of the NFS protocol, I find it useful to explicitly specify values for these parameters. This ensures that services that are only required by older NFS clients are also started, and also helps in debugging, because you know exactly what services are supposed to be started.**

Configuring Common NFS Options

Options for the common NFS services for an NFS server are set in the file `/etc/default/nfs-common`, which you can edit using your favorite text editor. Each of the following options is defined in a setting that looks like *`setting=value`*, where *`value`* is either `yes` or `no`. You should set the following options:

- `NEED_STATD`: Set this to `yes` to ensure that the `rpc.statd` daemon, required by NFS 2 and NFS 3 clients, is started. This server supports the NSM (Network Status Monitor) RPC protocol, which implements a reboot notification service (letting clients know when the NFS server has been rebooted), which is used by the NFS file locking service, `rpc.lockd`, to handle existing locks after a reboot. You can set this to `no` in a pure NFS 4 environment to avoid starting an unnecessary server, but I recommend against that unless you can guarantee that no client that speaks only the NFS 2 or NFS 3 protocol will ever need to mount directories from this NFS server.
- `STATDOPTS`: Enables you to pass arguments to the `rpc.statd` daemon when it starts. The most common of these are the `-p listen-port` and `-o send-port` options, which you can use to define specific ports for the status daemon to listen and send on. This can

be useful if you are using a port-based firewall or if you want to use ports other than the defaults as a security measure. Unless this is the case, you should usually not provide a value for this option.

- `NEED_IDMAPD`: Set this to `yes` to ensure that the `rpc.idmapd` daemon, required by NFS 4 clients, is started. This server supports UID/GID mappings to specific user and group names.
- `NEED_GSSD`: Set this to `no` to ensure that the `rpc.gssd` daemon, which is required for Kerberos authentication support, is not started. If you are using Kerberos authentication with NFS, you should set this to `yes` to ensure that this daemon is started.

After setting these values, save this file and restart the services associated with basic NFS support by executing the command `sudo /etc/init.d/nfs-common restart`.

NOTE **When run in start or restart modes, this startup script creates and mounts an in-memory pipe filesystem called `rpc_pipefs`, used for inter-process communication by RPC-based servers processes such as `rpc.idmapd` or `rpc.gssd`, mounting it on the directory `/var/lib/nfs/rpc_pipefs`. Although this filesystem will be listed in output from the mount command, you will not see this filesystem in your `/etc/fstab` file because the filesystem is created and mounted by the `nfs-common` script to limit NFS-related changes to files that are provided by other packages.**

Configuring NFS Kernel Server Options

Options for the common NFS services for an NFS server are set in the file `/etc/default/nfs-kernel-server`, which you can edit using your favorite text editor. Each of the following options is defined in a setting that looks like *`setting=value`*, where *`value`* is one of `yes` or `no`. The options that you should set are the following:

- `RPCNFSDCOUNT`: Set this to a specific numeric value to identify the number of NFS daemon instances that you want to start. The default value, 8, is usually sufficient, although you may want to increase this value if your NFS server is serving files to a large number of clients.
- `RPCNFSDPRIORITY`: Set this to a specific numeric value to start the NFS daemon at a specific priority. The numeric values to which you can set this option are the same as those used by the `nice` command. The default value, 0, is usually sufficient, although you may want to increase this value if your NFS server is serving files to a large number of clients and you want to ensure that it runs at a higher priority than other processes on the server.
- `RPCMOUNTDOPTS`: Enables you to pass arguments to the `rpc.mountd` daemon when it starts. The most common of these are the `-p listen-port` and `-o send-port` options, which you can use to define specific ports for the NFS mount daemon to listen and send on. This can be useful if you are using a port-based firewall or if you want to use ports other than the defaults as a security measure. Unless this is the case, you should usually not provide a value for this option.
- `NEED_SVCGSSD`: Set this to `no` to ensure that the `rpc.svcgssd` daemon, which is required for Kerberos authentication support, is not started. If you are using Kerberos authentication with NFS, you should set this to `yes` to ensure that this daemon is started.
- RPCSVCGSSDOPTS: Enables you to pass arguments to the `rpc.svcgssd` daemon when it starts. Most of the options for this daemon change the location or verbosity of its output, which you will only want to change if you are having authentication problems and want additional information about what it's doing internally. You should rarely need to provide a value for this option.

After setting these values, save this file and restart the services associated with the NFS server by executing the command `sudo /etc/init.d/nfs-kernel-server restart`.

NOTE **When run in start or restart modes, this startup script also creates and mounts an in-memory `nfsd` filesystem that is used for inter-process communication by instances of the NFS daemon, mounting it on the direct`/var/lib/nfs/rpc_pipefs`. Although this filesystem will be listed in output from the `mount` command, you will not see this filesystem in your `/etc/fstab` file because the filesystem is created and mounted by the `nfs-kernel-server` script to limit NFS-related changes to files that are provided by other packages.**

Setting Up the Pseudo-Filesystem and Exported Directories

As mentioned previously, NFS 4 differs from earlier versions of NFS by exporting a pseudo-filesystem in which the directories that you want to export are mounted using the mount command's `--bind` option. The pseudo-filesystem is uniquely identified by a special `fsid=0` parameter that you must specify in the `/etc/exports` options for that filesystem. Your `/etc/exports` file should contain entries for the pseudo-filesystem and for each directory that is mounted in it.

TIP **If you are exporting only a single directory, such as the mountpoint for a large shared filesystem, you can simply export that filesystem and use the `fsid=0` option to identify it as the top-level filesystem that you are exporting. However, this is relatively inflexible. If there is any chance that you will be exporting multiple directories/filesystems from an NFS server, it's smarter to set up your NFS server using a pseudo-filesystem from the beginning.**

The pseudo-filesystem used by an NFS 4 server is just a directory with special permissions, in which the directories that you want to export via NFS have been mounted. I typically use the directory `/export` on an NFS server as my pseudo-filesystem, but you can use any directory that you want. This directory should be created and its mode set using commands like the following:

```
sudo mkdir /export
sudo chmod a+rwxt /export
```

For each directory that you want to export via NFS, create a mountpoint for that directory under the directory that you just created and use the `mount` command's `--bind` option to mount the directory that you want to export there. For example, if you wanted to export the `/home` directory from your NFS server, you would execute the following commands:

```
sudo mkdir /export/home
sudo mount --bind /home /export/home
```

After repeating this process for each directory or mounted filesystem that you want to export, the next step in setting up your NFS server is to create appropriate entries for each of these in the file `/etc/exports`, as described in the next section.

Defining Exported Directories in /etc/exports

The file that contains information about each of the directories that is being exported by an NFS server is a text file named `/etc/exports`. Entries in this file have the following form:

```
full-path-name-of-exported-directory      hosts(export-options)
```

Each non-commented entry in the /etc/exports file is referred to as an export specification. The following are examples from one of my /etc/exports files:

```
/export       192.168.1.0/24(rw,fsid=0,insecure,no_subtree_check,async)
/export/users 192.168.1.0/24(rw,insecure,no_subtree_check,async)
```

The full-path-name-of-exported-directory field identifies the directory that is being exported. As shown in the previous example, your /etc/exports file should contain one export specification for your NFS pseudo-filesystem, and one entry for each directory that is mounted in that pseudo-filesystem.

The hosts entry identifies the hosts that can access a specific exported directory. Hosts can be specified by IP address, IP address and netmask, hostname, or subnet to state that only those hosts can access a specific directory exported by NFS. An entry such as 192.168.6.61 would limit access to a specific NFS directory from that host, while entries such as 192.168.6.*, 192.168.6.0, or 192.168.6.0/24 would limit access to a specific NFS directory to hosts on the 192.168.6 subnet. You can also grant access to an NFS server from any host that can access it by using an asterisk as the host specification.

The export-options identify specific characteristics of whether the exported directory is read-only or both readable and writable, how the UIDs and GIDs of incoming users should be mapped to users on the NFS server, how write and status requests should be processed, and so on. As you'd expect, many mount options are available. The pseudo-filesystem in which all of the directories that you are exporting are mounted must be identified using the fsid=0 export option, as shown in the previous example. The following are some of the more commonly used mount options:

- all_squash: Maps all NFS read or write requests to a specific user, usually "anonymous." This option is often used for public resources such as directories of USENET news, public FTP and download areas, and so on. All files written to an NFS directory that is exported with the all_squash mount option will be assigned the UID and GID of the user anonymous or some other UID and GID specified using the anonuid and anongid mount options. The default is no_all_squash, which preserves all UIDs and GIDs.
- async: Enables the NFS server to return control to the client before a write has completed, increasing performance. The default behavior when this option is not specified is synchronous writes, which are safer but which block the client until the write completes.
- insecure: Enables access to NFS directories by NFS clients that are running on nonstandard NFS network ports. By default, this option is off, and NFS requests must originate from ports where the port number is less than 1024. The insecure option may be necessary to enable access from random PC and Macintosh NFS clients. If you need to use this option, you should limit machines using the NFS option to a home network or secure corporate intranet. You should not use this option on any machines that are accessible from over the Internet because it introduces potential security problems.
- no_subtree_check: Prevents NFS from having to verify that a file requested through NFS is both in the filesystem from which the directory was exported on the NFS server and hierarchically under the directory that is being exported via NFS, known as a subtree check. Performing this check can cause problems when files in the NFS directory are renamed. This check is usually

disabled in filesystems with significant read/write activity by specifying the `no_subtree_check` mount option.

- `no_root_squash`: Lets root users on client workstations have the same privileges as the root user on the NFS file server. This option is off by default.
- `ro`: Exports directories that you don't want users to be able to write to because it is read-only. The default is `rw`, which enables read/write access.
- `sync`: Forces writes to the NFS server to be done synchronously, where the client waits for the writes to complete before returning control to the user. This is the default—as explained in the previous section, you can also specify asynchronous operation (`async`), which is slightly faster.

See the man page for `/etc/exports` (by using the `man 5 exports` command) for complete information on the options that are available in this file.

Once you have created an entry for a new exported directory in your `/etc/exports` file, you can export that directory by rerunning the `exportfs` command with the `-r` option, which tells the NFS server to reread the `/etc/exports` file and make any necessary changes to the list of directories that are exported by that NFS server.

Mounting NFS Directories on Clients

Many different export options are available for entries in an `/etc/exports` file export specification on an NFS server, but their number pales beside the number of mount options that you can specify when mounting an NFS filesystem from a client system. Mounting a directory from an NFS server can be done in several ways:

- By using the mount command from the command line, as in the following example, which mounts the entire pseudo-filesystem that is being exported from the NFS server `server` on the directory `/mnt`, using the mount options specified in `mount-options`. This pseudo-filesystem includes all of the entries that were mounted within it:

```
sudo mount -t nfs4 -o mount-options server:/ /mnt
```

- By using the mount command from the command line, as in the following example, which mounts the `/home` subdirectory that is being exported from the NFS server `server` on the directory `/mnt`, using the mount options specified in `mount-options`:

```
sudo mount -t nfs4 -o mount-options server:/home /mnt
```

- By creating an entry for that filesystem in the file `/etc/fstab`, and then mounting it by executing the mount command for that filesystem. The following is a sample `/etc/fstab` entry for the entire pseudo-filesystem that is being exported from the NFS server `server` on the directory `/mnt`, using the mount options specified in `mount-options`:

```
server:/   /mnt  nfs4  mount-options  0  0
```

TIP **Adding an entry to the `/etc/fstab` file is the best way to ensure that an NFS filesystem will always be mounted whenever the system is booted. You can maintain `/etc/fstab` entries for one or more NFS filesystems that you do not want to mount automatically by adding the `noauto` option to the *`mount-options`* string. Manually executing the mount command only mounts the filesystem until the system reboots or until you manually unmount it using the `umount` command. (There is no "n" after the initial "u" in `umount`, honest!)**

After executing any of the commands in this list, you should be able to access the specified NFS filesystem through the specified mount point. Access to these filesystems and their subdirectories will follow the authentication model that you are using: either standard UID/GID mappings or the authorization rights provided by an authentication model such as Kerberos.

One huge difference between mounting NFS 4 and NFS 2 or 3 exported directories is that NFS 4 filesystems are identified relative to the pseudo-filesystem that the NFS server is exporting, whereas NFS 2 or 3 filesystems required that you specify the full name of the directory that is being exported. For example, the NFS 2 or 3 equivalent of the third bullet in the previous list is the following:

```
sudo mount -t nfs -o mount-options server:/export/home /mnt
```

The filesystem type is `nfs`, rather than `nfs4`, and you must supply the full pathname of the exported filesystem. Your mount options specification should also include the string `nfsvers=3` to mount the exported directory in NFS 3 mode.

As mentioned previously, many mount options can be specified in a `mount-options` string when mounting NFS filesystems manually or in `/etc/fstab` entries. The `mount-options` string is a comma-separated sequence of mount options (and associated values, if required). The following are some of the most common:

- `hard`: Retry indefinitely if an NFS file operation times out. You cannot specify both `hard` and `soft` options.
- `intr`: Allows signals to interrupt NFS file operations that time out.
- `ro`: Mounts the NFS filesystem in read-only mode, regardless of how it is being exported by the NFS server.
- `rsize=`*`number`*: *`number`* specifies the size of the read buffer that the client should try to use when reading data from the NFS server.
- `rw`: Mounts the NFS filesystem in read-write mode if it being exported in read/write mode from the NFS server.
- `soft`: Reports an I/O error to the calling program when an NFS file operation times out. You cannot specify both the `soft` and `hard` options.
- `timeo=`*`number`*: *`number`* specifies the number of tenths of a second to wait before retransmitting when an RPC timeout occurs.
- `wsize=`*`number`*: *`number`* specifies the size of the write buffer that the client should try to use when writing back to the NFS server.

See the man page for the `nfs` command (by executing the command `man 5 nfs` from any command-line environment) for a complete list of mount options for NFS filesystems.

Automounting NFS Home Directories

Automounting is the process of automatically mounting NFS filesystems in response to requests for access to those filesystems. Automounting is controlled by an automount daemon that runs on the client system. In addition to automatically mounting filesystems in response to requests for access to them, an automount daemon can also automatically unmount volumes once they have not been used for a specified period of time.

Using an automount daemon prevents you from having to mount shared NFS directories that you are not actually using at the moment. Mounting all NFS directories on all clients at all times causes a reasonable amount of network traffic, much of which is extraneous if those directories are not actually being used. Using the NFS automount daemon helps keep NFS-related network traffic to a minimum.

The automount daemon is called `autofs` and is available in the autofs package. The `autofs` automounter daemon is configured through the file `/etc/auto.master` or through NIS maps with the same name. Because the `autofs` daemon is part of the kernel, it is faster and the automounter of choice for many Linux distributions (such as Ubuntu). Both of these packages are available in the Ubuntu repositories, but discussing all of the nuances of automounting is outside the scope of this chapter.

Debugging NFS Problems

If you encounter problems mounting an NFS filesystem or accessing the files that it contains, there are a number of ways to try to resolve the problem. The next few sections provide suggestions to help you identify and resolve the problem.

Verifying RPC Services

If you are having problems mounting an NFS filesystem, the first thing to check is to use the `rpcinfo` command to probe the NFS server that you are trying to contact to verify that it is running both the right NFS services and the right versions of those services. The `-p` option probes the specified host, or the current system if no host name or IP address is provided. The following example probes the `portmapper` on the NFS server at the IP address 192.168.6.143:

```
$ rpcinfo -p 192.168.6.143
   program vers proto   port
    100000    2   tcp    111  portmapper
    100000    2   udp    111  portmapper
    100024    1   udp  33264  status
    100024    1   tcp  37575  status
    100003    2   udp   2049  nfs
    100003    3   udp   2049  nfs
    100003    4   udp   2049  nfs
    100021    1   udp  56344  nlockmgr
    100021    3   udp  56344  nlockmgr
    100021    4   udp  56344  nlockmgr
    100003    2   tcp   2049  nfs
```

```
100003    3   tcp   2049  nfs
100003    4   tcp   2049  nfs
100021    1   tcp  50947  nlockmgr
100021    3   tcp  50947  nlockmgr
100021    4   tcp  50947  nlockmgr
100005    1   udp  54674  mountd
100005    1   tcp  33465  mountd
100005    2   udp  54674  mountd
100005    2   tcp  33465  mountd
100005    3   udp  54674  mountd
100005    3   tcp  33465  mountd
100005    1   tcp  33465  mountd
100005    2   udp  54674  mountd
100005    2   tcp  33465  mountd
100005    3   udp  54674  mountd
100005    3   tcp  33465  mountd
```

This output shows that NFS versions 3 and 4 are both supported by the NFS server and the lock manager. In this case, the right services seem to be available on the server. You should also verify that the required processes are running on both the NFS server and the client system by using the `ps` command with your favorite options for listing system processes, as in the following example from an NFS 4 server:

```
$ ps -ef | egrep '(rpc|nfs)'
root   16447   2  0 Jul20 ?      00:00:00 [rpciod/0]
statd  17308   1  0 Jul20 ?      00:00:00 /sbin/rpc.statd
root   17316   1  0 Jul20 ?      00:00:00 /usr/sbin/rpc.idmapd
root   17345   2  0 Jul20 ?      00:00:00 [nfsd4]
root   17346   2  0 Jul20 ?      00:00:00 [nfsd]
root   17347   2  0 Jul20 ?      00:00:00 [nfsd]
root   17348   2  0 Jul20 ?      00:00:00 [nfsd]
root   17349   2  0 Jul20 ?      00:00:00 [nfsd]
root   17350   2  0 Jul20 ?      00:00:00 [nfsd]
root   17351   2  0 Jul20 ?      00:00:00 [nfsd]
root   17352   2  0 Jul20 ?      00:00:00 [nfsd]
root   17353   2  0 Jul20 ?      00:00:00 [nfsd]
root   17357   1  0 Jul20 ?      00:00:00 /usr/sbin/rpc.mountd
root   17415   2  0 Jul20 ?      00:00:00 [nfsv4-svc]
root   17828 945  0 05:47 tty1   00:00:00 egrep (rpc|nfs)
```

An NFS 4 server should be running at least the `rpc.idmapd`, `rpc.mountd`, and `nfsd` processes. An NFS 4 client should be running at least the `rpc.idmapd` process.

Verifying Filesystem Export

The `kernel-nfs-server` package includes utilities called `exportfs` and `showmount` that you can use to list the directories that an NFS server is currently exporting from your system and verify that the correct filesystems are being exported. The following is an example of using the `exportfs` command's `-v` to display a verbose listing of the filesystems that are being exported:

```
$ exportfs -v
/export/users  192.168.6.0/24(rw,async,wdelay,nohide,\
     insecure,no_root_squash,no_subtree_check)
/export/pub    192.168.6.0/24(rw,async,wdelay,\
```

```
     insecure,root_squash,all_squash,no_subtree_check)
/export        192.168.6.0/24(rw,async,wdelay,insecure,\
     root_squash,no_subtree_check,fsid=0)
```

(Entries for each exported directory would appear on a single line in actual output.)

The following is an equivalent example of using the `showmount` command:

```
$ showmount -e
Export list for userver:
/export       192.168.6.0/24
/export/pub   192.168.6.0/24
/export/users 192.168.6.0/24
```

If the right processes are running and the right filesystems are being exported, it may be necessary to bring out the big guns to help diagnose your problem—activating RPC debugging, as described in the next section.

For complete information about the `exportfs` and `showmount` utilities, see their online reference information, which is available by typing `man exportfs` or `man showmount` from any Ubuntu command line, such as an xterm or the GNOME Terminal application.

Debugging NFS Operations

If you are having problems mounting or accessing NFS directories, you can generate tracing information in the system log (`/var/log/messages`) by activating various RPC debugging options. The easiest way to activate these is to write a specific debugging mask to the files in the `/proc` filesystem that are associated with these debugging options. These files are the following:

- `/proc/sys/sunrpc/nfs_debug`: Controls tracing for the NFS client. The tracing output is generated in the system log on the host where the NFS client is running—the host that is attempting to mount or access an NFS directory. The debugging values that are available are listed in `/usr/include/linux/nfs_fs.h`.
- `/proc/sys/sunrpc/nfsd_debug`: Controls tracing for the NFS daemon. The tracing output is generated in the system log on the host where the NFS daemon is running—the host that is exporting the NFS directory. The debugging values that are available are listed in `/usr/include/linux/nfsd/debug.h`.
- `/proc/sys/sunrpc/nlm_debug`: Controls tracing for the Network Lock manager. The tracing output is generated in the system log on the host where the NFS Lock Manager is running—the host that is exporting the NFS directory. The debugging values that are available are listed in `/usr/include/linux/lockd/debug.h`.
- `/proc/sys/sunrpc/rpc_debug`: Controls tracing for the underlying RPC layer used by all NFS client/server communications. The debugging values that are available are listed in `/usr/include/linux/sunrpc/debug.h`.

You can write decimal numbers that represent a bitmask that enables different types of tracing to these files by using the `sysctl` command or by simply writing to the appropriate file. For example, the following two commands activate all possible debugging for the NFS client:

```
# echo 65535 > /proc/sys/sunrpc/nfs_debug
# sysctl -w sunrpc.nfs_debug=65535
```

If you can think in terms of bitmasks, you can consult the appropriate kernel header files from the previous list. Personally, I typically find the NFS client and NFS server output to be most useful, and typically begin with a value of 1023 because this activates most of the useful debugging for each without truly becoming overwhelming. I often test both the client and the server on the server by first starting the server, and then attempting to mount one of its exported directories somewhere else on the same host. Once things work locally, you can move on to try to get them working from a remote client.

Getting More Information About NFS

Not surprisingly, the Web provides many excellent sources of additional information about NFS. For more information, consult any of the following:

- NFS FAQ (`http://nfs.sourceforge.net`)
- NFS HOWTO (`http://nfs.sourceforge.net/nfs-howto`)
- NFS4 on Ubuntu (`http://help.ubuntu.com/community/NFSv4Howto`)

Summary

This chapter explained NFS, the Network File System, which is the de facto standard for sharing directories between UNIX-like systems over a network. NFS is simple, lightweight, and fast—and its implementation has been freely available for almost any UNIX-like system since the mid 1980s. After a bit of background, this chapter explained how to install and configure an NFS server on your Ubuntu desktop, server, or Kubuntu system(s).

Although NFS implementations for Microsoft Windows systems exist, most of them cost money, and thus they aren't an out-of-the-box alternative for people who have heterogeneous networks comprised of systems running Linux, Mac OS X, and Microsoft Windows. (Mac OS X systems, with their Mach and BSD underpinnings, natively support NFS.) Interestingly enough, the most "standard" way of sharing files between all of these types of systems is by using a package known as Samba, which allows Ubuntu and other UNIX-like systems to share directories and filesystems over a network using the proprietary Server Message Block (SMB) protocol used by Microsoft Windows systems (and from which Samba takes its name). Bonus Chapter 14 explains how to quickly and easily set up Samba so that everyone on your home network can access files stored on your Ubuntu and Kubuntu systems, regardless of what operating system they are running, and (of course) at no cost.

Bonus Chapter 14

Setting Up a Samba Server

Like it or not, the planet is infested with Windows machines. As you can see from statements like that, I'm probably as guilty as anyone of propagating the "us vs. them" mentality when it comes to Windows vs. Linux. How I personally feel about Windows and Microsoft really doesn't matter—the important thing in this chapter is to discuss the various ways in which software available on Linux systems makes it easy to integrate Linux and Windows filesystems in both directions, and providing features such as automatic printer sharing as freebies along the way.

You'd have to have been living in a cave for the last five or so years not to have heard of Samba, arguably one of the most popular applications ever written for Linux and UNIX-like systems. In a nutshell, Samba is a set of applications that was originally developed to provide support for Microsoft's networking protocols on Linux systems, but which has been ported to just about every other network-aware operating system.

A huge number of books are available that are dedicated to discussing Samba, explaining every nuance of its configuration files, installation, and use. My goal in this chapter is not to embed another one inside a book on Ubuntu, but rather to provide some interesting background information about Windows networking and Samba, and then to explain how to use Samba to share directories and printers on your Ubuntu system so that Microsoft Windows users in your home, academic, or business computing environment can access them.

IN THIS CHAPTER

Installing a Samba server

Sharing directories with Windows systems

Sharing printers with Windows systems

Overview of Microsoft Windows File Sharing

Networking and related technologies such as routing are probably responsible for more acronyms than any other aspect of the computer industry. MS-DOS and Microsoft Windows networking have contributed their share, largely

because of the ubiquity of these operating systems in modern computing environments. Because of the popularity of DOS systems (yesterday) and Windows systems (today), today's Windows systems provide support for almost everyone's networking protocols. Frankly, Windows does an admirable job of continuing to make forward progress while still maintaining backward compatibility with almost every ancient DOS application and networking protocol. Windows systems still support the Internet Packet Exchange (IPX) and Sequenced Packet Exchange (SPX) networking protocols used by Novell to provide the first PC file servers. However, more relevant for our discussion here are the networking protocols and attendant acronyms that were developed by Microsoft and used to provide file and resource sharing over PC networks without requiring the involvement of any third parties, thank you very much.

The *Basic Input and Output System* used by PCs to interact with local devices is best known by its initials as the PC's *BIOS*. As networks began to appear, Microsoft extended the capabilities of the BIOS to support accessing and sharing information over a network, naming the related protocols the network BIOS, or as it's more popularly known, *NetBIOS*. Just as the BIOS provides the basic functions that support all system input and output, the NetBIOS provides the basic functions that let you use and administer network services. NetBIOS commands and functions must be exchanged between networked systems and therefore require a lower-level network transport mechanism to move network packets from one host to another. The lower-level transport protocols that are still in common use in PC networking today are IPX, NetBEUI (Network Basic Extended User Interface), and TCP/IP (Transmission Control Protocol/Internet Protocol). Interestingly, the "Internet" in the full names of both IPX and TCP/IP refers to inter-network communications, not the Internet as we know it today.

Modern Windows systems send their NetBIOS requests by using TCP/IP as a transport protocol. On top of the NetBIOS level, Windows networking provides a higher-level interface for network services known as the Server Message Block (SMB) protocol, which is a networking protocol that can be easily used by applications. SMB is a connection-oriented protocol rather than a broadcast protocol, meaning that it depends on establishing connections to specific networked services provided by other networked hosts rather than simply broadcasting its availability. Once a connection is established, SMB provides four basic types of functions:

- Session functions that negotiate and establish networked connections between machines (often referred to as virtual circuits), authenticate, and verify the access privileges that each party has with the other.
- File functions that enable applications to open, close, read, and write remote files, shared directories, and so on.
- Printer functions that enable applications to spool output to remote output devices.
- Message functions that enable applications to send and receive control, status, and informational messages between different systems on the network.

SMB became an Open Group standard for networking interoperability in the early 1990s. Samba takes its name from SMB, the addition of two vowels making it easily pronounced and somewhat softer than simply being YAA (Yet Another Acronym).

In recent developments, an enhanced version of the SMB protocol called CIFS (Common Internet File System) was submitted by Microsoft to the IETF (Internet Engineering Task Force), an open association of people that are interested in the architecture of Internet communication and the smooth operation of the Internet. CIFS has been approved as a standard, and extends the capabilities of SMB by expanding its focus to sharing resources using even more open, cross-platform standards such as HTTP URLs (HyperText Transfer Protocol Uniform Resource Locators) and DNS (the Domain Name System used to map hostnames to IP addresses and vice versa).

Introducing Samba

When you get right down to it, more data is probably stored on Windows systems than on any other type of computer system. All of those 10GB home and office systems add up to a tremendous number of Windows filesystems holding a staggering amount of data. Samba gives Linux users transparent access to Windows filesystems, but is more commonly used to give Windows users transparent access to Linux, UNIX, and UNIX-like systems. Samba does this by providing a network interface that is compatible with the networked file and printer-sharing protocols used between Windows systems. To a Windows system, a Linux system running Samba looks exactly like a random Windows system that is sharing filesystems across the network. This enables Windows users to take advantage of the speed, power, and capacity of Linux systems without even realizing that they are accessing Linux filesystems.

Samba is a free and impressive interface for Linux, UNIX, and other types of systems to any other networked device that can communicate using the SMB protocol, most notably Windows systems that provide networked access to files, directories, and printers. Samba enables Windows users to access Linux file systems and resources just like any other Windows shared file system or networked resource. For example, with Samba running on a Linux system on your network, Windows users can mount their Linux home directories as networked Windows drives and automatically print to Linux printers just like any other networking Windows printer. Samba, which was originally authored by Andrew Tridgell, is one of the most impressive pieces of interoperability software ever developed.

Samba includes both client and server software—in other words, client software that enables users to communicate from Linux machines to SMB hosts on your network, and server software that provides an SMB interface for your Linux machine. Using the Samba client software is discussed in Bonus Chapter 2. This chapter focuses on explaining how to install and set up a Samba server.

A Samba server actually consists of two processes, both of which can be started from the command line or automatically by integrating them into your system's startup procedure. These processes are `smbd`, the Samba daemon that provides file sharing and print services to Windows clients, and `nmbd`, the NetBIOS name server that maps the NetBIOS names used by Windows SMB requests to the IP addresses used by Linux systems. The Samba daemon is configured by modifying its configuration file, `/etc/samba/smb.conf`. On Ubuntu and Kubuntu systems, the easiest and most standard way to configure specific directories that you want to export via Samba is to manually modify the Samba configuration file, as discussed later in the section "Samba Server Configuration Essentials."

Interoperability between Linux and Windows systems requires much more than just Samba. The Linux kernel provides built-in support for the protocols used to access Windows filesystems, enabling Linux users to mount Windows filesystems via entries in `/etc/fstab`, just like any other filesystem resource.

NOTE **The Ubuntu repositories provide both the current version of Samba and preliminary packages for the next version of Samba, Samba 4. This chapter focuses on the current version of Samba (Samba 3.*x*) because it is completely stable. Although Samba 4 will deliver improvements in performance, power, and flexibility, it isn't ready for prime time use at the moment. Stay tuned!**

Installing the Samba Server and Friends

Samba client packages are installed by default on Ubuntu and Kubuntu desktop systems, making it easy for any Ubuntu or Kubuntu system to contact and access SMB/CIFS shared directories from Microsoft Windows systems or Samba servers. To install the packages required to run and monitor a Samba server so that you can share directories from your Ubuntu or Kubuntu system, you must install the `samba`

package using your favorite package management tool: Synaptic, Adept, `aptitude`, or `apt-get`. See Chapter 23 for more detailed information on installing using these package management tools.

When installing the samba package, you may also want to install the `samba-doc` and `samba-doc-pdf` packages, which respectively provide HTML and PDF versions of all of the official Samba project documentation, plus an online copy of a book entitled *Samba 3 By Example* by one of the leaders of the Samba project, John Terpstra.

Samba Server Configuration Essentials

While graphical utilities for configuring Samba exist, they are in a state of flux on Ubuntu and Kubuntu systems. Previous releases of Ubuntu included a Shared Folders tool that enabled you to identify the Linux directories that you wanted to export via NFS or Samba, but that tool is no longer part of Ubuntu. A tool called system-config-samba, originally developed for the Red Hat Linux distribution, is available for Ubuntu systems, but doesn't provide access to all of the configuration options that you will want to use in most cases. As a graphical tool, it also cannot be used on Ubuntu server systems unless you also install a window manager or desktop environment. Similarly, a KDE control center plug-in was available for KDE-based systems such as Kubuntu, but hasn't been updated to work within the new System Settings framework. Therefore, this chapter focuses on doing Samba server configuration in the aging but tried-and-true Linux way—by editing configuration files using a text editor. Samba's configuration file is `/etc/samba/smb.conf`. The Samba configuration file contains many helpful comments, which are lines beginning with a hash mark. It also contains many sample, inactive configuration commands, which are lines beginning with a semicolon. These indicate configuration commands that you may want to activate by removing the leading semicolon.

As you'll see in the next few sections, editing the file `/etc/samba/smb.conf` to configure Samba is actually quite simple, but is also somewhat inelegant when compared to the rest of the system administrative environment provided by Ubuntu Linux.

NOTE **Web-based system administration tools such as `swat` (Samba Web Administration Tool) and `webmin` (a more generalized Web-based administrative environment) are available from the Ubuntu repositories, but require special configuration before they'll work correctly in the Ubuntu or Kubuntu environments. Rather than taking the conceptual detour of explaining how to use these tools, I'll continue to wait for a desktop solution for Samba setup and configuration on desktop Ubuntu and Kubuntu systems.**

The next few sections focus on specific entries in the Samba configuration file that you may need to update to use Samba in your environment, instead of discussing the Samba configuration file in its entirety. Between the official Samba documentation, the online documentation that is provided as part of the `samba-doc` and `samba-doc-pdf` packages, and the many books available on Samba, there's no shortage of detailed information about every nuance of Samba—the rest of this chapter focuses on getting things to work, as quickly and as easily as possible.

Identifying Your Workgroup or Domain

The Samba configuration file is divided into several sections, each identified by a section name enclosed within square brackets. The entries related to the network identity of your Samba server and how you authenticate to it are located in the `[global]` section of the Samba configuration file `/etc/samba/smb.conf`.

The key entry that you must set to define how your Ubuntu server interacts with Windows systems is the workgroup entry, which identifies the workgroup or Windows domain to which your Samba server belongs. On a sample system of mine, this entry looks like this:

```
[global]
   workgroup = WVH
```

This entry identifies my machine as belonging to the Windows workgroup or domain named WVH. In this SOHO example, it's a workgroup, but that is essentially transparent—the key thing here is that your Samba server is either a member of the workgroup/domain that you are already using at your site, or a primary domain controller that defines the domain that you want it to host.

NOTE **If you want your Samba server to function as the primary domain controller (PDC) for a Windows domain, you must also change the value of the domain master setting in the [Domains] section of the Samba configuration file to yes, and make sure that the entry is not commented out. In this case, the entry would look something like the following:**

```
domain master = auto
```

This means that the Samba server will serve as a primary domain controller for your domain if no other domain controller can be located, which will be the case if your Samba server is hosting the domain.

This chapter explains how to use Samba as part of a workgroup, which is the typical way that Samba is used on a home or SOHO network. For information about setting up a complete Windows domain, make sure that you installed the `samba-doc` as suggested in the section of this chapter "Installing the Samba Server and Friends," and consult the online copy of *Samba 3 By Example* that is provided in the directory `/usr/share/doc/samba-doc/htmldocs/Samba3-ByExample`.

Because this example shows a small workgroup, you will probably also want to activate the entries that tell Samba not to function as a WINS (the Windows NetBIOS name service, discussed earlier in this chapter), but to instead use your system's Domain Name Server to look up Windows hostnames, which should look like the following:

```
wins support = no
dns proxy = yes
name resolve order = lmhosts host wins bcast
```

The last entry tells Samba that clients will first look in their local `lmhosts` file for name information (on Windows systems), will then check information in the file `/etc/hosts` on the Samba server, will then check WINS (which, in this case, is proxied to the local DNS server), and will finally use a broadcast to search for the right host. This combination covers the hostname lookup bases for every Samba server configuration that I've ever set up or used.

Finally, you must define the security model that will be used to identify authenticated users of the Samba server. To do this, uncomment the following entry in the Samba configuration file:

```
security = user
```

This completes the core modifications to the Samba configuration file that are necessary for Samba to be able to share files with the specified workgroup or domain. That wasn't so bad, was it?

Configuring Samba Authentication

In the workgroup configuration that we are using as an example in this chapter, Samba comes preconfigured to perform the type of authentication required by workgroup members. However, to be able to access shared resources, you must be able to authenticate to the Samba server. Samba maintains its own authentication information—to add login and password information about a user, you must use the `smbpasswd` command. When adding information about a user, you must use the `-a` option, followed by the name of the user that you want to add. The `smbpasswd` command will prompt you for a password and will then prompt you for that password again to ensure that you have typed the password correctly. The following is a sample transcript of adding the user `wvh`:

```
$ sudo smbpasswd -a wvh
Password:
New SMB password:
Retype new SMB password:
```

The first password prompt is from the `sudo` application, asking for my Ubuntu password so that I can perform this privileged operation. The second two are for entering and verifying the Samba password for the user `wvh`.

NOTE **The suggested security level in the Samba configuration file, `security = user`, requires that Samba users must also exist in the password file on the Ubuntu system, or you will receive an error message and the `smbpasswd` command will fail. For information about other security models and their implications, see the file `/usr/share/doc/samba-doc/htmldocs/ServerType.html` from the `samba-doc` package.**

Sharing Printers and Home Directories Using Samba

After setting up the workgroup and creating a user, the next thing to consider is the general resources on the Ubuntu system that you want to make available to all users of Windows systems that connect to your Samba server. The most common examples of these are printers and users' home directories.

The entries in the `[global]` section of a Samba configuration file that are relevant to printers and printing on Ubuntu are the following, which you should make sure are not commented out (i.e., preceded by a semicolon):

```
load printers = yes
printing = cups
printcap name = cups
```

Later in the Samba configuration file, the `[printers]` and `[print$]` sections provide information about how your Windows system will interact with the Samba server. You won't need to change any of these, but for your reference, these entries are as follows:

```
[printers]
   comment = All Printers
   browseable = no
   path = /tmp
   printable = yes
   guest ok = no
   read only = yes
   create mode = 0700
 [print$]
```

```
   comment = Printer Drivers
   path = /var/lib/samba/printers
   browseable = yes
   read only = yes
   guest ok = no
```

The `[printers]` section identifies how the Samba server will handle requests for printer identification and incoming print requests from Windows clients, while the `[print$]` section maps the traditional Windows `print$` share name to a directory on your Linux system where you can put the print drivers for specific printers so that your Windows clients can locate and load them if they are not already available on a client system.

Aside from printers, the most natural resource for users to want to access from their Windows systems is their Ubuntu or Kubuntu home directory. This provides an easy way to use your home directory on a Samba server as a centralized place to store files, but also enables you to automatically back up files from your Windows system to the Samba server and enables you to leverage any standard backups that you are doing on the Samba server. (See Bonus Chapter 6 for more information about backing up files on your Ubuntu system.)

Happily, Samba's configuration file is already set up to support sharing home directories by default, as long as your Linux and Windows user names match. The following is the section of the Samba configuration file that supports this:

```
[homes]
   comment = Home Directories
   browseable = no
   read only = yes
   create mask = 0700
   directory mask = 0700
```

As you can see from this configuration file excerpt, home directories are shared by default but are not writable from Windows clients. This is pretty inconvenient, so you will want to change the `read only` entry in this section of the configuration file to `no`, as in the following example:

```
read only = no
```

This enables users to both read and write files in their home directories.

TIP **If your Windows and Ubuntu users happen to have different logins, you can associate the two by creating appropriate entries in the file `/etc/samba/smbusers`. Entries in this file have the form `UnixLogin = WindowsLogin`. For example, to map the Ubuntu user `wvh` to the Windows user `bill.vonhagen`, I would create the following entry in the `/etc/samba/smbusers` file:**

```
wvh = bill.vonhagen
```

If your Windows login names contain one or more spaces, you must enclose them within quotation marks in this file.

In practice, I suggest keeping logins the same across Ubuntu and Samba systems, but this file can be used if you have different naming conventions for the users of different types of systems.

Verifying the Samba Configuration File

After making any changes to your Samba configuration file, you should verify that you haven't accidentally violated the syntax of the file. If a Samba configuration file contains any invalid entries, Samba displays an error message and will not load. This could be extremely discouraging, were it not for the fact that the Samba server package also provides a utility that tests the validity of a Samba configuration file and identifies the exact location of any errors that it finds. This utility is `testparm`, which you can run with no arguments to test the default Samba configuration file `/etc/samba/smb.conf`. The output from a run of the `testparm` utility looks like the following:

```
$ testparm
Load smb config files from /etc/samba/smb.conf
Processing section "[homes]"
Processing section "[printers]"
Processing section "[print$]"
Loaded services file OK.
WARNING: passdb expand explicit = yes is deprecated
Server role: ROLE_STANDALONE
Press enter to see a dump of your service definitions
```

As shown in the preceding example, after displaying some general information, the `testparm` utility prompts you to press Enter on your keyboard to see detailed information about the services that your Samba server has been configured to provide. The following is an example of this detailed information:

```
[global]
        workgroup = WVH
        server string = %h server (Samba, Ubuntu)
        obey pam restrictions = Yes
        passdb backend = tdbsam, guest
        syslog = 0
        log file = /var/log/samba/log.%m
        max log size = 1000
        name resolve order = lmhosts host wins bcast
        printcap name = cups
        panic action = /usr/share/samba/panic-action %d
        invalid users = root
        printing = cups
        print command =
        lpq command = %p
        lprm command =

[homes]
        comment = Home Directories
        read only = No
        create mask = 0700
        directory mask = 0700
        browseable = No

[printers]
        comment = All Printers
        path = /tmp
        create mask = 0700
```

```
        printable = Yes
        browseable = No

[print$]
        comment = Printer Drivers
        path = /var/lib/samba/printers
```

Depending on the version of Samba installed on your Ubuntu system and the changes that you have made to your Samba configuration file, your output may differ slightly.

The `testparm` utility only checks the Samba configuration file—it does not consult the running Samba service for information about how it is configured and the services that it provides. To do this, you can use the `smbclient` utility, as described in the next section.

Testing Samba Availability and Services

The Samba applications suite includes an application called `smbclient`, which is a command-line client for contacting and browsing Samba servers. You can also use this utility to test if your Samba server is up and running, and to list the services that it is providing.

If you have made any changes to your Samba configuration file, you should restart your Samba server before attempting to verify the services that it provides. To do this, execute the following command:

```
/etc/init.d/samba restart
```

This command shuts down any instance of the Samba server that is currently running, and then starts the Samba server again, which forces it to read your updated configuration file.

Once you have done this, you can query your Samba server by using the `smbclient` application, as discussed in Bonus Chapter 2, to verify that your Samba server is working. Once it's clear that the Samba server is working and is correctly configured to allow users to access user-specific resources, it's time to add some system-wide resources that all authenticated users will be able to take advantage of. Luckily, as explained in the next section, you can even do this with a graphical tool.

Sharing Other Directories via Samba

Samba's configuration file makes it easy to access your home directory on a Samba server because that's probably the directory that you will want to access most frequently from a Windows system. As mentioned earlier, being able to access you Linux home directory from a Windows system provides a convenient mechanism for backing up critical Window files by copying them to your Linux home directory (which you're certainly backing up, right?). Of course, if you are using a Samba server as a domain controller and using Samba's network logon mechanism to actually use Samba directories as your Windows home directories, things are even easier.

At some point, however, you will certainly want to export other directories from your Samba server so that all users can have access to them. For example, I always share the `/opt2` directory that contains my online audio collection, which I should point out was either legally ripped from CDs that I own or consists of freely downloadable live recordings. To do this, I add the following entries to the end of my `/etc/samba/smb.conf` file:

```
[opt2]
  path = /opt2
```

```
comment = Music
available = yes
browseable = no
public = yes
read only = yes
```

These entries enable all users on my network to browse for and access this share. Adding these entries to the standard Samba configuration file ensures that the newly defined share `opt2` will always be available as soon as the Samba daemon is started (or restarted) on the Ubuntu or Kubuntu system that is acting as a Samba server.

Getting More Information About Samba

Not surprisingly, the Web is an excellent source of additional information about Samba. Somewhat surprisingly, the Samba project itself provides some excellent documentation, as well as an online version of a great book about the current version of Samba, which is Samba version 3. As suggested earlier, you can install your own copies of the Samba documentation by installing the `samba-doc` and `samba-doc-pdf` packages when you install the Samba server. For more information, consult any of the following:

- `www.samba.org`: The main site for the Samba project, where you can find the latest tips, tricks, and source code.
- `/usr/share/doc/samba-doc/htmldocs/Samba3-HOWTO`: A directory containing a HOWTO file in HTML format that explains how to install and configure Samba version 3, and answers many common questions. (You'll want to open the file `index.html` in this directory from your Web browser.) The HTML version of this document is provided in the `samba-doc` package.
- `/usr/share/doc/samba-doc/htmldocs/Samba3-ByExample`: A directory containing an HTML version of an excellent, hands-on book about Samba 3. This book was written by John Terpstra, one of the leaders of the Samba project, and explains how to install and configure Samba 3 for a complete spectrum of networked environments, from smaller SOHO environments (as discussed in this chapter) to enterprise environments with thousands of users. (You'll want to open the file `index.html` in this directory from your Web browser.) The HTML version of this document is provided in the `samba-doc` package. You'll probably also want to buy a paper copy of John's book—it's definitive, complete, and battery-free.

Summary

This chapter explained how to install and set up a Samba server to enable your Ubuntu or Kubuntu system to export directories so that they can be accessed by users who are still stuck on Microsoft Windows systems. The combination of Samba utilities and the SMB/CIFS support that is built into the Linux kernel also provides the ability to access Windows shares from Ubuntu or Kubuntu systems, all of which is available by default on those systems, as explained in Bonus Chapter 2.

Samba is probably the best example of the value that open source software can bring to modern computing, enabling seamless data sharing between UNIX-like and Microsoft Windows systems thanks to a fantastic combination of reverse engineering and insightful development.

Everything you need to know about Ubuntu—and then some!

It's the hottest Linux distribution out there and now you can join the excitement, with Ubuntu 8.10 and this information-packed guide. Want to view Flash files? Filter incoming mail? Rip audio CDs? Connect to Windows, FTP servers, or other systems? Learn all this and much more with the expert guidance, tips, and techniques in this complete Ubuntu reference. From the basics for newcomers to enterprise management for system administrators, it's what you need to succeed with Ubuntu.

- Master Ubuntu fundamentals for desktop and networks
- Send e-mail, share files, edit text, and print
- Store personal files in private, encrypted folders
- Explore the Gnome Desktop, add widgets, and create themes
- Work with the Bash shell and the command line
- Set up Web, mail, print, Samba, NFS, and other servers
- Create virtual machines, go wireless, and set up Wi-Fi



What's on the DVD and CD-ROM?



On the DVD, you'll find the ready-to-use download of Ubuntu 8.10 Linux.



On the CD, fourteen bonus chapters discuss advanced networking, software development, system policies and authentication, administration tasks, virtual machines, and setting up servers, as well as a complete pdf of the book.

System Requirements: Please see the "About the DVD and CD-ROM Appendix" for details and complete system requirements.

www.wiley.com/compbooks

WILEY
wiley.com

Shelving Category:
COMPUTERS/Operating Systems/Linux
Reader Level:
Beginning to Advanced

$49.99 USA
$59.99 Canada

ISBN 978-0-470-29420-8
54999
9 780470 294208